
Can Watermarking Large Language Models Prevent Copyrighted Text Generation and Hide Training Data?

Michael-Andrei Panaitescu-Liess¹ Zora Che¹ Bang An¹ Yuancheng Xu¹ Pankayaraj Pathmanathan¹
Souradip Chakraborty¹ Sicheng Zhu¹ Tom Goldstein¹ Furong Huang¹

Abstract

Large Language Models (LLMs) have demonstrated impressive capabilities in generating diverse and contextually rich text. However, concerns regarding copyright infringement arise as LLMs may inadvertently produce copyrighted material. In this paper, we first investigate the effectiveness of watermarking LLMs as a deterrent against the generation of copyrighted texts. Through theoretical analysis and empirical evaluation, we demonstrate that incorporating watermarks into LLMs significantly reduces the likelihood of generating copyrighted content, thereby addressing a critical concern in the deployment of LLMs. Additionally, we explore the impact of watermarking on Membership Inference Attacks (MIAs), which aim to discern whether a sample was part of the pretraining dataset and can be used to detect copyright violations. Surprisingly, we find that watermarking adversely affects the success rate of MIAs, complicating the task of detecting copyrighted text in the pretraining dataset. Finally, we propose an adaptive technique to improve the success rate of a recently developed MIA under watermarking. Our findings underscore the importance of developing adaptive methods to study critical problems in LLMs with potential legal implications.

1. Introduction

In recent years, Large Language Models (LLMs) have pushed the frontiers of natural language processing by facilitating sophisticated tasks like text generation, translation, and summarization. With their impressive performance, LLMs are increasingly integrated into various applications, including virtual assistants, chatbots, content generation,

¹University of Maryland, College Park, USA. Correspondence to: Michael-Andrei Panaitescu-Liess <mpanaite@umd.edu>.

Work presented at the Next Generation of AI Safety Workshop at ICML 2024, Vienna, Austria. Copyright 2024 by the author(s).

and education. However, the widespread usage of LLMs brings forth serious concerns regarding potential copyright infringements. Addressing these challenges is critical for the ethical and legal deployment of LLMs.

Copyright infringement involves unauthorized usage of copyrighted content, which violates the intellectual property rights of copyright owners, potentially undermining content creators' ability to fund their work, and affecting the diversity of creative outputs in society. Additionally, violators can face legal consequences, including lawsuits and financial penalties. For LLMs, copyright infringement can occur through (1) generation of copyrighted content during deployment and (2) illegal usage of copyrighted works during training. Ensuring the absence of copyrighted content in the vast training datasets of LLMs is challenging. Moreover, legal debates around generative AI copyright infringement vary by region, complicating compliance further.

Current lawsuits against AI companies for unauthorized use of copyrighted content (e.g., *Andersen v. Stability AI Ltd*, *NYT v. OpenAI*) highlight the urgent need for methods to address these challenges. In this paper, we focus on watermarking LLMs to tackle two main issues: (1) preventing the generation of copyrighted content, and (2) detecting copyrighted content in training data. We show that watermarking can significantly impact both the generation of copyrighted text and the detection of copyrighted content in training set.

Firstly, we observe that current LLM output watermarking techniques can significantly reduce the probability of generating copyrighted content, by tens of orders of magnitude. Our empirical results focus on two recent watermarking methods: UMD (Kirchenbauer et al., 2023) and Unigram-Watermark (Zhao et al., 2023). Both methods split the vocabulary into two sets (green and red) and bias the model towards selecting tokens from the green set by altering the logits distribution, thereby embedding a detectable signal. We provide both empirical and theoretical results to support our findings (see Section 3 and Appendix A.4 for details).

Secondly, we demonstrate that watermarking techniques can decrease the success rate of Membership Inference Attacks (MIAs), which aim to detect whether a piece of copyrighted




Figure 1. Illustration of the effect of LLM watermarking on generation of copyrighted content. We observe that watermarking can reduce the probability of generating copyrighted content by more than 10^{20} times on Llama-30B.

text was part of the training dataset. Since MIAs exploit the model’s output, their performance can suffer under watermarking due to changes in the probability distribution of output tokens. Our comprehensive empirical study, including 5 recent MIAs and 5 LLMs, shows that the AUC of detection methods can be reduced by up to 16% in the presence of watermarks.

Finally, we propose an adaptive method designed to enhance the success rate of a recent MIA (Shi et al., 2023) under watermarking, to improve detection of copyright violation under watermarking. This method applies a correction to the model’s output to account for the perturbations introduced by watermarks. By incorporating knowledge about the watermarking scheme, we improve the detection performance for pretraining data, counteracting the obfuscation caused by watermarking. Our contribution underscores the importance of continuously developing adaptive attack methodologies to keep pace with advances in defense mechanisms.

The rest of the paper is organized as follows. We formally introduce the problems that we study in Section 2 and present the empirical results for the first two contributions in Sections 3 and 4. Due to space limitations, we include our theoretical analysis in Appendix A.4 and we introduce the adaptive version of the Min-K% Prob membership inference attack in Appendix A.5. Finally, in Section 5, we discuss the limitations of our work and provide concluding remarks. We discuss prior work on LLM watermarking, copyright, memorization, and membership inference and we present additional experiments in Appendix A.1, A.2, A.3 and A.6.

2. Setup and Notations

2.1. Definitions

Let D be a training dataset, let C be the set of all the copyrighted texts, and let C_D be all the copyrighted texts that are part of D . We give definitions for the following setups.

Verbatim Memorization of Copyrighted Content. Let P be the set of all possible prompts. For a language model f , trained on D and a piece of copyrighted text $c \in C$, we define c as being memorized verbatim if $\max_{p \in P} \mathbb{P}(c = f(p)) > t$, where t is a threshold that needs to be set. In other words, a piece of text is memorized verbatim if it is generated as the output of some prompt with high

probability. Instead of the intractable computation of \max over all prompts, we measure memorization using a continuous and tractable metric: the perplexity of the model on the copyrighted text c when given a fixed prompt p . In our experiments, we consider p as an empty string or the set of the first 10, 20 or 100 tokens of c . Lower perplexity values thereby indicate higher levels of verbatim memorization.

MIAs for Copyrighted Training Data Detection. MIAs are privacy attacks aiming to detect whether a sample was part of the training set. We define an MIA for copyrighted data as a binary classifier $A(\cdot)$, which ideally outputs $A(x) = 1, \forall x \in C_D$ and $0, \forall x \in C - C_D$. In practice, $A(\cdot)$ is defined by thresholding a metric (e.g., perplexity), i.e., $A(x) = 1, \forall x$ such that $\text{perplexity}(x) < t$ and 0 , otherwise. Since the threshold t needs to be set, prior work (Shi et al., 2023) uses AUC (Area Under the ROC Curve) as an evaluation metric which is independent of t . Note that we employ the same metric in our experiments.

LLM Watermarking. Watermarking LLMs consists of introducing signals during its training or inference that are difficult to detect by humans without the knowledge of a *watermark key* but can be detected using an algorithm if the key is known. We focus our paper on recent methods that employ logits distribution changes as a way of inserting watermark signals during the decoding process (Kirchenbauer et al., 2023) (Zhao et al., 2023).

2.2. Membership Inference Attacks (MIAs)

Current MIAs for detecting training data rely on thresholding various heuristics that capture differences in output probabilities for each token between data included in the training set and data that was not. Below, we present an overview of these heuristics.

Perplexity. This metric represents the perplexity of the LLM on a given sample, as discussed in the previous section.

Smaller Ref, Lowercase and Zlib (Carlini et al., 2021). These heuristics are defined as follows:

- **Smaller Ref:** The ratio of the log-perplexity of the target LLM on a sample to the log-perplexity of a smaller reference LLM on the same sample.
- **Lowercase:** The ratio of the log-perplexity of the target

LLM on the original sample to the log-perplexity of the LLM on the lowercase version of the sample.

- Zlib: The ratio of the log-perplexity of the target LLM on a sample to the zlib entropy of the sample.

Min-K% Prob (Shi et al., 2023). This heuristic computes the average of the minimum $K\%$ token probabilities outputted by the LLM on the sample. Note that this method requires tuning K , so in all our experiments we chose the best result over $K\% \in \{5\%, 10\%, 20\%, 30\%, 40\%, 50\%, 60\%\}$.

2.3. LLM Watermarking Methods

UMD (Kirchenbauer et al., 2023) splits the vocabulary into two sets (green and red) and biases the model towards the green tokens by altering the logits distribution. The hash of the previous token’s ID serves as a seed for a pseudo-random number generator used to split the vocabulary into these two groups. For a “hard” watermark, the model is forced not to sample from the red list at all. For a “soft” watermark, a bias δ is added to the logits of the green tokens before sampling. We focus our empirical evaluation on “soft” watermarks as they are more suitable for LLM deployment due to their smaller impact on the quality of the generated text.

Unigram-Watermark (Zhao et al., 2023) employs a similar approach of splitting the vocabulary into two sets and biasing the model towards one of the two sets. However, the split remains consistent throughout the generation of tokens. This choice is made to provide a provable improvement against paraphrasing attacks (Krishna et al., 2024).

3. Watermarking LLMs Prevents Copyrighted Text Generation

In this section, we study the effect of watermarking methods on verbatim memorization. We discuss the implications of watermarking for preventing copyright text generation.

Datasets. We consider 4 versions of the WikiMIA benchmark (Shi et al., 2023) with 32, 64, 128, and 256 words in each sample and only consider the samples that were very likely part of the training set of all the models we consider (labeled as 1 in Shi et al. (2023)). We consider these subsets as a proxy for text that was used in the training set, and the model may be prone to verbatim memorization. From now on, we refer to this subsets as the “training samples” or “training texts”. Similarly, we consider BookMIA dataset (Shi et al., 2023), which contains samples from copyrighted books, in Appendix A.6.

Metric. We measure the relative increase in perplexity on the watermarked model’s training samples compared to the original model. We report an increase in the minimum and

average perplexity over the training samples. Note that a large increase in perplexity corresponds to a large decrease in the probability of generating that specific sample, as shown later in this section. When computing the perplexity, we prompt the model with an empty string, the first 10, and the first 20 tokens of the targeted training sample, respectively. For BookMIA dataset we also consider the first 100 tokens as a prompt because the number of words in each sample is 512, which is larger than in the case of WikiMIA.

Models. We conduct our empirical evaluation using 5 recent LLMs: Llama-30B, Llama-13B (Touvron et al., 2023), GPT-NeoX-20B (Black et al., 2022), Pythia-2.8B (Biderman et al., 2023) and OPT-2.7B (Zhang et al., 2022).

3.1. Empirical Evaluation

Table 1. Measuring the reduction in verbatim memorization of training texts on WikiMIA-32. We report the relative increase in perplexity between the watermarked and unwatermarked models, where larger values correspond to less memorization. Note that “P:” stands for “prompt length”.

	P.	Llama-30B		NeoX-20B		Llama-13B	
		Min.	Avg.	Min.	Avg.	Min.	Avg.
UMD	0	3.3	31.2	3.7	52.1	4.9	34.3
	10	2.8	28.7	2.2	52.1	3.5	31.9
	20	2.4	30.1	1.8	66.0	3.5	33.4
Unigram	0	4.1	34.1	4.4	54.1	5.0	36.6
	10	3.0	31.7	2.8	52.5	4.0	34.3
	20	2.4	31.5	2.0	56.4	3.4	34.0

In Table 1, we show the increase in perplexity on the training samples when the model is watermarked relative to the unwatermarked model. We observe that for Llama-30B, Unigram-Watermark induces a relative increase of 4.1 in the minimum and 34.1 in the average perplexity. Note that the former represents a decrease of at least 4.3×10^{22} in the probability of generating the sample (and this is in the case the sample has only 32 tokens, which is likely a lower bound since the number of tokens is normally larger than the number of words). We observe consistent results over several models and prompt lengths. Note that we use a fixed strength parameter $\delta = 10$ for the watermark methods in these experiments and a fixed percentage of green tokens, 50%, which we use in all the experiments from this paper unless specified otherwise. All the results are averaged over 5 runs using different seeds for the watermark methods. We include additional results on WikiMIA-64, WikiMIA-128 and WikiMIA-256 in Tables 4, 5 and 6, respectively, in Appendix A.2. We observe that our findings are consistent across models and splits of WikiMIA.

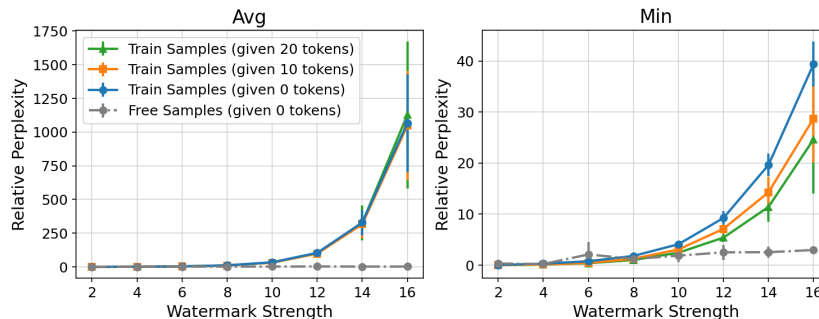


Figure 2. We study the impact of watermark strength on the relative increase in average and minimum perplexity of the training samples on WikiMIA-32 and the quality of the generated text. “Free samples” is the baseline of generating text freely. There is an exponential increase in the perplexity of the training texts as watermark strength increases, while free generation quality is affected at a slower rate.

In Figure 2, we study the influence of the watermark strength δ on the relative increase in both minimum and average perplexity on the training samples from WikiMIA-32. In this experiment, we also consider a baseline of generating text freely to study the impact of watermarks on the quality of text relative to the impact on training samples’ generation (here, perplexity is computed by an unwatermarked model). All the results are averaged over 5 runs using different seeds for the watermark methods. In the case of free generation, we generate 100 samples for 5 different watermarking seeds and average the results. The length of the generated samples is up to 42 tokens, which is approximately 32 words in the benchmark (on a token-to-word ratio of 4 : 3). **The results show an exponential increase in the perplexity of the training texts with the increase in watermark strength, while the generation quality is affected at a slower rate.** This suggests that even if there is a trade-off between protecting the generation of text memorized verbatim and generating high-quality text, **finding a suitable watermark strength for each particular application is possible.**

Takeaways. Watermarking significantly increases the perplexity of generating training texts, reducing verbatim memorization likelihood. This is achieved with only a moderate impact on the overall quality of the generated text, suggesting that effective watermark strength can be tailored to balance verbatim memorization and text quality for specific applications. Finally, we believe that our findings directly extend to the generation of copyrighted text verbatim, as this constitutes a form of verbatim memorization of the training data. Since copyrighted texts are not expected to be distributed significantly differently from the rest of the training data, the probability of generating copyrighted materials under watermarking is also likely to decrease. To confirm, we run experiments on a dataset containing copyrighted data (BookMIA) and include the results in the Appendix A.6.

4. Impact of Watermarking on Pretraining Data Detection

Datasets. We consider the WikiMIA benchmark again as in Section 3. We consider the full datasets, rather than the subset of samples that were part of the training for models we are considering. Also, we consider BookMIA dataset containing copyrighted texts in Appendix A.6.

Metrics. We follow the prior work (Shi et al., 2023; Duarte et al., 2024) and report the AUC and AUC drop to study the detection performance of the MIAs. Note that this metric has the advantage of not having to tune the threshold for the detection classifier.

Models. We conduct experiments on the same 5 LLMs as in Section 3. Additionally, for the Smaller Ref method that requires a smaller reference model along with the target LLM, we consider Llama-7B, Neo-125M, Pythia-70M, and OPT-350M as references.

4.1. Empirical evaluation

In Table 2, we show the AUC for the unwatermarked and watermarked models using the UMD scheme, as well as the drop between the two. We observe that watermarking reduces the AUC (drop shown in bold in the table) by up to 14.2% across 4 detection methods and 5 LLMs. Note that all the experiments on watermarked models are run with 5 different seeds and we report the mean and standard deviation of the results. We also report the AUC drop, which is computed by the difference between the AUC for the unwatermarked model and the mean AUC over the 5 runs for the watermarked model. Additionally, while the experiments from Table 2 are conducted on WikiMIA-256, we observe similar trends for WikiMIA-32, WikiMIA-64, and WikiMIA-256 in Appendix A.3. We also study the impact of the watermark’s strength on the AUC drop for Llama-30B and GPT-Neo-20B in Figure 3 and for the other models in




Figure 3. We study the AUC drop due to watermarking for each MIA methods when varying the strength of the watermark. We observe that higher watermark strengths generally induce larger AUC drops.

Figure 4 from the Appendix A.3 (we consider WikiMIA-256 for these experiments). We observe that higher watermark strengths generally induce larger AUC drops.

In addition to the 4 detection methods, we also consider Smaller Ref attack, which we include in Table 11 of Appendix A.3. We consider different variations, including an unwatermarked reference model and a watermarked one with a similar strength but a different seed or with both strength and seed changed in comparison to the watermarked target model. The baseline is an unwatermarked model with an unwatermarked reference model. We observe the AUC drops in all scenarios (up to 16.4%), which is consistent with our previous findings.

Finally, we experiment with several percentages of green tokens for a fixed watermark strength of $\delta = 10$. We show the results in Table 12 of Appendix A.3. We observe that for all models, in at least 80% of the cases all of the attacks’ AUCs are negatively affected (positive drop value), suggesting that **finding a watermarking scheme that reduces the success rates of the current MIAs is not difficult**. Note that the experiments are run on WikiMIA for UMD scheme and the results are averaged over 5 watermark seeds.

Takeaways. Watermarking can significantly reduce the success of membership inference attacks (MIAs), with AUC drops up to 16.4%. By varying the percentage of green tokens as well as the watermark’s strength, we observe that watermarking schemes can be easily tuned to negatively impact the detection success rates of MIAs.

5. Conclusion and Discussion

Watermarking LLMs has unintended consequences on methods towards copyright protection. Our experiments demonstrate that while watermarking may be a promising solution to prevent copyrighted text generation, watermarking also complicates membership inference attacks that may be employed to detect copyright abuses. Watermarking can be a

Table 2. AUC of MIAs for the unwatermarked (*top* of each cell), watermarked (UMD scheme) models (*middle* of each cell), and the drop between the two (*bottom* of each cell) on WikiMIA-256.

	Llama-30B	NeoX-20B	Llama-13B
PPL	72.0%	71.3%	71.2%
	$70.6 \pm 1.9\%$ 1.4%	$64.7 \pm 2.3\%$ 6.6%	$70.0 \pm 2.6\%$ 1.2%
Lowercase	68.1%	68.2%	65.5%
	$63.8 \pm 4.5\%$ 4.3%	$55.4 \pm 5.5\%$ 14.2%	$61.6 \pm 3.8\%$ 3.9%
Zlib	72.7%	73.2%	73.1%
	$72.0 \pm 1.6\%$ 0.7%	$66.6 \pm 2.0\%$ 6.6%	$71.6 \pm 2.3\%$ 1.5%
Min-K% Prob	71.8%	78.0%	72.9%
	$70.5 \pm 1.8\%$ 1.3%	$76.2 \pm 2.1\%$ 1.8%	$70.4 \pm 3.2\%$ 2.5%

double-edged sword for copyright regulators since it promotes compliance during generation time, while making training time copyright violations harder to detect. We hope our work further the discussion around watermarking and copyright issues for LLMs.

Limitations & Future Work. Our work considers only decoding time watermarking techniques, future work may benefit from studying other types of watermarking methods, and explore whether watermarked models can generate paraphrased versions of copyrighted text as this is another copyright infringement concern. Our proposed method for improving MIA success rate on watermarked model makes strong assumption on the watermarking scheme, which may not always be satisfied despite empirical improvements in our experiments. Our observations on the deterioration of MIA success suggests that for copyright violation auditing, an unwatermarked model or the watermarking scheme may be needed. We encourage the community to further refine adaptive methods to ensure robust copyright protection and data privacy, and consider the interactions of different methods on downstream legal concerns.

Acknowledgements

Panaitescu-Liess, Che, An, Xu, Pathmanathan, Chakraborty, Zhu, and Huang are supported by DARPA Transfer from Imprecise and Abstract Models to Autonomous Technologies (TIAMAT) 80321, National Science Foundation NSF-IIS-2147276 FAI, DOD-ONR-Office of Naval Research under award number N00014-22-1-2335, DOD-AFOSR-Air Force Office of Scientific Research under award number FA9550-23-1-0048, DOD-DARPA-Defense Advanced Research Projects Agency Guaranteeing AI Robustness against Deception (GARD) HR00112020007, Adobe, Capital One and JP Morgan faculty fellowships.

This work was made possible by the ONR MURI program and the AFOSR MURI program. Commercial support was provided by Capital One Bank, the Amazon Research Award program, and Open Philanthropy. Further support was provided by the National Science Foundation (IIS-2212182), and by the NSF TRAILS Institute (2229885).

Impact Statement

Protecting copyrighted materials and detecting copyright violations are important societal issues as large language models are rapidly adopted in daily life. Our work demonstrates that watermarking has unintended effects on methods towards copyright protection. We believe our work has implications for both model builders and model auditors. Our results suggest that watermarking may be a double-edged sword, decreasing the chance of copyright violation at generation time, yet also making copyright abuse during training time harder to detect.

References

Aaronson, S. Simons institute talk on watermarking of large language models. 2023. URL <https://simons.berkeley.edu/talks/scott-aaronson-ut-austin-openai-2023-08-17>

Bentley, J. W., Gibney, D., Hoppenworth, G., and Jha, S. K. Quantifying membership inference vulnerability via generalization gap and other model metrics. *arXiv preprint arXiv:2009.05669*, 2020.

Biderman, S., Schoelkopf, H., Anthony, Q. G., Bradley, H., O’Brien, K., Hallahan, E., Khan, M. A., Purohit, S., Prashanth, U. S., Raff, E., et al. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pp. 2397–2430. PMLR, 2023.

Black, S., Biderman, S., Hallahan, E., Anthony, Q., Gao, L., Golding, L., He, H., Leahy, C., McDonnell, K., Phang,

J., et al. Gpt-neox-20b: An open-source autoregressive language model. *arXiv preprint arXiv:2204.06745*, 2022.

Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., and Song, D. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *28th USENIX security symposium (USENIX security 19)*, pp. 267–284, 2019.

Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U., et al. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pp. 2633–2650, 2021.

Carlini, N., Chien, S., Nasr, M., Song, S., Terzis, A., and Tramer, F. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 1897–1914. IEEE, 2022a.

Carlini, N., Ippolito, D., Jagielski, M., Lee, K., Tramer, F., and Zhang, C. Quantifying memorization across neural language models. *arXiv preprint arXiv:2202.07646*, 2022b.

Carlini, N., Hayes, J., Nasr, M., Jagielski, M., Schwag, V., Tramer, F., Balle, B., Ippolito, D., and Wallace, E. Extracting training data from diffusion models. In *32nd USENIX Security Symposium (USENIX Security 23)*, pp. 5253–5270, 2023.

Christ, M., Gunn, S., and Zamir, O. Undetectable watermarks for language models. *arXiv preprint arXiv:2306.09194*, 2023.

Chu, T., Song, Z., and Yang, C. How to protect copyright data in optimization of large language models? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 17871–17879, 2024.

Duarte, A. V., Zhao, X., Oliveira, A. L., and Li, L. De-cop: Detecting copyrighted content in language models training data. *arXiv preprint arXiv:2402.09910*, 2024.

Elkin-Koren, N., Hacohen, U., Livni, R., and Moran, S. Can copyright be reduced to privacy? *arXiv preprint arXiv:2305.14822*, 2023.

Hacohen, U., Haviv, A., Sarfaty, S., Friedman, B., Elkin-Koren, N., Livni, R., and Bermano, A. H. Not all similarities are created equal: Leveraging data-driven biases to inform genai copyright disputes. *arXiv preprint arXiv:2403.17691*, 2024.

Ippolito, D., Tramèr, F., Nasr, M., Zhang, C., Jagielski, M., Lee, K., Choquette-Choo, C. A., and Carlini, N. Preventing generation of verbatim memorization in language models gives a false sense of privacy. In *Proceedings*

- of the 16th International Natural Language Generation Conference, pp. 28–53. Association for Computational Linguistics, 2023.
- Kandpal, N., Wallace, E., and Raffel, C. Deduplicating training data mitigates privacy risks in language models. In *International Conference on Machine Learning*, pp. 10697–10707. PMLR, 2022.
- Karamolegkou, A., Li, J., Zhou, L., and Søggaard, A. Copyright violations and large language models. *arXiv preprint arXiv:2310.13771*, 2023.
- Kirchenbauer, J., Geiping, J., Wen, Y., Katz, J., Miers, I., and Goldstein, T. A watermark for large language models. In *International Conference on Machine Learning*, pp. 17061–17084. PMLR, 2023.
- Krishna, K., Song, Y., Karpinska, M., Wieting, J., and Iyyer, M. Paraphrasing evades detectors of ai-generated text, but retrieval is an effective defense. *Advances in Neural Information Processing Systems*, 36, 2024.
- Kuditipudi, R., Thackstun, J., Hashimoto, T., and Liang, P. Robust distortion-free watermarks for language models. *arXiv preprint arXiv:2307.15593*, 2023.
- Lee, K., Ippolito, D., Nystrom, A., Zhang, C., Eck, D., Callison-Burch, C., and Carlini, N. Deduplicating training data makes language models better. *arXiv preprint arXiv:2107.06499*, 2021.
- Liu, A., Pan, L., Hu, X., Meng, S., and Wen, L. A semantic invariant robust watermark for large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=6p8lpe4MNf>.
- Mattern, J., Mireshghallah, F., Jin, Z., Schoelkopf, B., Sachan, M., and Berg-Kirkpatrick, T. Membership inference attacks against language models via neighbourhood comparison. In Rogers, A., Boyd-Graber, J., and Okazaki, N. (eds.), *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 11330–11343, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.719. URL <https://aclanthology.org/2023.findings-acl.719>.
- Mireshghallah, F., Goyal, K., Uniyal, A., Berg-Kirkpatrick, T., and Shokri, R. Quantifying privacy risks of masked language models using membership inference attacks. *arXiv preprint arXiv:2203.03929*, 2022.
- Oren, Y., Meister, N., Chatterji, N., Ladhak, F., and Hashimoto, T. B. Proving test set contamination in black box language models. *arXiv preprint arXiv:2310.17623*, 2023.
- Ren, J., Xu, H., He, P., Cui, Y., Zeng, S., Zhang, J., Wen, H., Ding, J., Liu, H., Chang, Y., et al. Copyright protection in generative ai: A technical perspective. *arXiv preprint arXiv:2402.02333*, 2024.
- Sablayrolles, A., Douze, M., Schmid, C., Ollivier, Y., and Jégou, H. White-box vs black-box: Bayes optimal strategies for membership inference. In *International Conference on Machine Learning*, pp. 5558–5567. PMLR, 2019.
- Shejwalkar, V., Inan, H. A., Houmansadr, A., and Sim, R. Membership inference attacks against nlp classification models. In *NeurIPS 2021 Workshop Privacy in Machine Learning*, 2021.
- Shi, W., Ajith, A., Xia, M., Huang, Y., Liu, D., Blevins, T., Chen, D., and Zettlemoyer, L. Detecting pretraining data from large language models. *arXiv preprint arXiv:2310.16789*, 2023.
- Shokri, R., Stronati, M., Song, C., and Shmatikov, V. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pp. 3–18. IEEE, 2017.
- Somepalli, G., Singla, V., Goldblum, M., Geiping, J., and Goldstein, T. Diffusion art or digital forgery? investigating data replication in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6048–6058, 2023a.
- Somepalli, G., Singla, V., Goldblum, M., Geiping, J., and Goldstein, T. Understanding and mitigating copying in diffusion models. *Advances in Neural Information Processing Systems*, 36:47783–47803, 2023b.
- Song, C. and Shmatikov, V. Auditing data provenance in text-generation models. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 196–206, 2019.
- Tirumala, K., Markosyan, A., Zettlemoyer, L., and Aghajanyan, A. Memorization without overfitting: Analyzing the training dynamics of large language models. *Advances in Neural Information Processing Systems*, 35: 38274–38290, 2022.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., and Lample, G. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Vyas, N., Kakade, S. M., and Barak, B. On provable copyright protection for generative models. In *International Conference on Machine Learning*, pp. 35277–35299. PMLR, 2023.

Wen, Y., Liu, Y., Chen, C., and Lyu, L. Detecting, explaining, and mitigating memorization in diffusion models. In The Twelfth International Conference on Learning Representations, 2024. URL <https://openreview.net/forum?id=84n3UwkH7b>.

Yeom, S., Giacomelli, I., Fredrikson, M., and Jha, S. Privacy risk in machine learning: Analyzing the connection to overfitting. In 2018 IEEE 31st computer security foundations symposium (CSF), pp. 268–282. IEEE, 2018.

Zhang, J., Sun, J., Yeats, E., Ouyang, Y., Kuo, M., Zhang, J., Yang, H., and Li, H. Min-k%++: Improved baseline for detecting pre-training data from large language models. arXiv preprint arXiv:2404.02936, 2024.

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

Zhao, X., Ananth, P., Li, L., and Wang, Y.-X. Provable robust watermarking for ai-generated text. arXiv preprint arXiv:2306.17439, 2023.

A. Appendix

A.1. Related work

Watermarks for LLMs. Language model watermarking techniques embed identifiable markers into output text to detect AI-generated content. Recent strategies incorporate watermarks during the decoding phase of language models (Zhao et al., 2023; Kirchenbauer et al., 2023). Aaronson (2023) develops the Gumbel watermark, which employs traceable pseudo-random sampling for generating subsequent tokens. Kirchenbauer et al. (2023) splits the vocabulary into red and green lists according to preceding tokens, biasing the generation towards green tokens. Zhao et al. (2023) employs a fixed grouping strategy to develop a robust watermark with theoretical guarantees. Liu et al. (2024) proposes to generate watermark logits based on the preceding tokens’ semantics rather than their token IDs to boost the robustness of the watermark. Kuditipudi et al. (2023) and Christ et al. (2023) explore watermark methods that do not change the output textual distribution.

Copyright. Copyright protection in the age of AI has gained importance, as discussed by Ren et al. (2024). Vyas et al. (2023) addresses content protection through near access-freeness (NAF) and developed learning algorithms for generative models to ensure compliance under NAF conditions. Prior works focus on training algorithms to prevent copyrighted text generation (Vyas et al., 2023; Chu et al., 2024), whereas our work emphasizes lightweight, inference-time algorithms. Other works have studied copyright in machine learning from a legal perspective. Hacothen et al. (2024) utilizes a generative model to determine the generic characteristics of works to aid in defining the scope of copyright. Elkin-Koren et al. (2023) demonstrates that copying does not necessarily constitute copyright infringement and argues that existing detection methods may detract from the foundational purposes of copyright law.

Memorization. One cause of copyright issues is that machine learning models may memorize training data. Prior studies have observed that LLMs can memorize copyrighted or private information in training data, such as phone numbers and addresses (Karamolegkou et al., 2023; Carlini et al., 2019; 2021; Lee et al., 2021), leading to significant privacy and security concerns. To measure memorization, Carlini et al. (2021) proposes eidetic memorization, defining a string as memorized if it was present in the training data and it can be reproduced by a prompt. This definition, along with variations like exact and perfect memorization, has been widely adopted in subsequent studies (Tirumala et al., 2022; Kandpal et al., 2022). Carlini et al. (2022b) quantitatively measures memorization in LLMs as the fraction of extractable training data and finds that memorization significantly grows as model size scales and training examples are duplicated. To minimize memorization, Lee et al. (2021) and Kandpal et al. (2022) propose deduplicating training data, which also improves accuracy. Ippolito et al. (2023) proposes an inference time defense that perfectly prevents all verbatim memorization. However, it can not prevent the leakage of training data due to the existence of many “style-transfer” prompts, suggesting it is a challenging open problem. Unlike the methods that we are studying in this paper, Ippolito et al. (2023) requires access to a complete set of copyrighted texts that the model was trained on. Memorization in the image domain has also been studied from various angles (Somepalli et al., 2023a;b; Carlini et al., 2023; Wen et al., 2024).

Membership Inference. As a proxy for measuring memorization, membership inference attacks (MIAs) predict whether or not a particular example was used to train the model (Shokri et al., 2017; Yeom et al., 2018; Bentley et al., 2020). Most membership inference attacks rely only on the model’s loss since the model is more likely to overfit an example if it is in the training data (Sablayrolles et al., 2019). Carlini et al. (2022a) trains shadow models to predict whether an example is from the training data. In the NLP domain, many works have focused on masked language models (Miresghallah et al., 2022) and fine-tuning data detection (Song & Shmatikov, 2019; Shejwalkar et al., 2021). Recently, Shi et al. (2023) studies pretraining data inference and introduced a detection method based on the hypothesis that unseen examples are likely to contain outlier words with low probabilities under the LLM. Zhang et al. (2024) approaches pretraining data detection by measuring how sharply peaked the likelihood is around the inputs. Duarte et al. (2024) proposes detecting copyrighted content in training data by probing the LLM with multiple-choice questions, whose options include both verbatim text and their paraphrases. Other methods include testing perplexity differences (Mattern et al., 2023) and providing provable guarantees of test set contamination without access to pretraining data or model weights (Oren et al., 2023).

A.2. Additional experiments on verbatim memorization

Table 3. Measuring the reduction in verbatim memorization of training texts on WikiMIA-32. We report the relative increase in perplexity between the watermarked and unwatermarked models, where larger values correspond to less memorization. Note that “P” stands for “prompt length”.

		Llama-30B		NeoX-20B		Llama-13B		Pythia-2.8B		OPT-2.7B		
		P.	Min.	Avg.	Min.	Avg.	Min.	Avg.	Min.	Avg.	Min.	Avg.
UMD	0	3.3	31.2	3.7	52.1	4.9	34.3	11.4	61.3	10.4	64.5	
	10	2.8	28.7	2.2	52.1	3.5	31.9	8.8	63.7	8.3	67.7	
	20	2.4	30.1	1.8	66.0	3.5	33.4	5.0	74.0	7.0	84.4	
Unigram	0	4.1	34.1	4.4	54.1	5.0	36.6	14.3	74.5	11.5	66.1	
	10	3.0	31.7	2.8	52.5	4.0	34.3	11.8	73.6	9.8	70.2	
	20	2.4	31.5	2.0	56.4	3.4	34.0	6.6	79.1	5.8	81.4	

Table 4. Measuring the reduction in verbatim memorization of copyrighted texts on WikiMIA-64. We report the relative increase in perplexity between the watermarked and unwatermarked model, so larger values correspond to less memorization. Note that "P." stands for "prompt length".

		Llama-30B		NeoX-20B		Llama-13B		Pythia-2.8B		OPT-2.7B		
		P.	Min.	Avg.	Min.	Avg.	Min.	Avg.	Min.	Avg.	Min.	Avg.
UMD	0	4.9	27.6	4.2	42.8	6.7	30.5	15.2	50.4	14.9	51.7	
	10	4.3	26.2	3.7	41.3	6.1	29.1	15.6	49.2	14.4	51.4	
	20	3.9	26.4	3.6	43.1	5.8	29.3	14.0	50.3	12.5	52.7	
Unigram	0	5.0	28.1	4.3	45.3	6.7	30.9	17.6	62.3	16.0	53.7	
	10	3.8	26.9	3.4	43.6	5.3	29.7	16.2	60.6	17.1	53.0	
	20	3.2	26.9	3.1	44.2	4.4	29.7	13.6	60.9	11.7	53.6	

Table 5. Measuring the reduction in verbatim memorization of copyrighted texts on WikiMIA-128. We report the relative increase in perplexity between the watermarked and unwatermarked model, so larger values correspond to less memorization. Note that "P." stands for "prompt length".

		Llama-30B		NeoX-20B		Llama-13B		Pythia-2.8B		OPT-2.7B		
		P.	Min.	Avg.	Min.	Avg.	Min.	Avg.	Min.	Avg.	Min.	Avg.
UMD	0	5.7	25.3	4.6	39.5	7.6	28.0	23.1	45.3	18.6	48.1	
	10	5.3	24.4	4.3	38.9	7.2	27.1	23.6	44.7	19.1	47.6	
	20	5.2	24.5	4.3	39.3	6.8	27.2	23.0	44.7	17.5	47.8	
Unigram	0	4.5	25.6	5.9	42.9	6.4	28.2	17.6	54.9	19.6	50.0	
	10	3.9	25.0	5.3	42.0	5.7	27.6	15.8	53.6	18.7	49.9	
	20	3.6	25.2	5.1	42.1	5.3	27.7	15.1	53.6	18.0	49.9	

Can Watermarking Large Language Models Prevent Copyrighted Text Generation and Hide Training Data?

Table 6. Measuring the reduction in verbatim memorization of copyrighted texts on WikiMIA-256. We report the relative increase in perplexity between the watermarked and unwatermarked model, so larger values correspond to less memorization. Note that "P." stands for "prompt length".

		Llama-30B		NeoX-20B		Llama-13B		Pythia-2.8B		OPT-2.7B		
		P.	Min.	Avg.	Min.	Avg.	Min.	Avg.	Min.	Avg.	Min.	Avg.
UMD	0	7.5	23.9	15.4	37.8	13.0	26.3	31.2	45.4	27.3	46.3	
	10	7.3	23.4	15.5	37.6	12.5	25.8	30.9	45.2	27.5	46.0	
	20	7.2	23.5	16.1	37.6	12.6	25.9	30.4	45.0	27.6	46.2	
Unigram	0	7.4	24.4	21.0	42.4	13.9	26.8	36.9	54.3	28.8	46.3	
	10	7.1	24.1	21.2	41.9	13.7	26.5	35.4	53.7	28.2	46.0	
	20	6.7	24.2	21.5	41.8	13.7	26.5	34.6	53.4	29.3	45.9	

A.3. Additional experiments on pretraining data detection

Table 7. AUC of each MIA for the unwatermarked (*top* of each cell), watermarked models (*middle* of each cell), and the drop between the two (*bottom* of each cell) on WikiMIA-256 using UMD scheme.

	Llama-30B	NeoX-20B	Llama-13B	Pythia-2.8B	OPT-2.7B
PPL	72.0%	71.3%	71.2%	67.8%	60.5%
	70.6 ± 1.9%	64.7 ± 2.3%	70.0 ± 2.6%	64.4 ± 1.9%	54.9 ± 2.2%
	1.4%	6.6%	1.2%	3.4%	5.6%
Lowercase	68.1%	68.2%	65.5%	62.9%	58.9%
	63.8 ± 4.5%	55.4 ± 5.5%	61.6 ± 3.8%	58.7 ± 3.2%	49.7 ± 2.9%
	4.3%	14.2%	3.9%	4.2%	9.2%
Zlib	72.7%	73.2%	73.1%	69.2%	62.7%
	72.0 ± 1.6%	66.6 ± 2.0%	71.6 ± 2.3%	66.1 ± 1.2%	58.1 ± 1.8%
	0.7%	6.6%	1.5%	3.1%	4.6%
Min-K% Prob	71.8%	78.0%	72.9%	71.0%	65.5%
	70.5 ± 1.8%	76.2 ± 2.1%	70.4 ± 3.2%	69.5 ± 1.6%	63.1 ± 3.4%
	1.3%	1.8%	2.5%	1.5%	2.4%

Table 8. AUC of each MIA for the unwatermarked (*top* of each cell), watermarked models (*middle* of each cell) and the drop between the two (*bottom* of each cell) on WikiMIA-128 using UMD scheme.

	Llama-30B	NeoX-20B	Llama-13B	Pythia-2.8B	OPT-2.7B
PPL	70.3%	70.6%	67.7%	62.8%	60.0%
	66.3 ± 2.2%	63.6 ± 2.4%	63.4 ± 2.6%	61.4 ± 2.3%	55.1 ± 1.6%
	4.0%	7.0%	4.3%	1.4%	4.9%
Lowercase	59.1%	68.0%	60.6%	59.4%	57.1%
	55.9 ± 2.9%	58.2 ± 3.4%	55.1 ± 3.0%	55.7 ± 1.6%	49.2 ± 4.5%
	3.2%	9.2%	5.5%	3.7%	7.9%
Zlib	71.8%	72.3%	69.6%	64.9%	62.3%
	68.6 ± 2.3%	66.3 ± 2.1%	65.8 ± 2.7%	63.9 ± 1.9%	58.9 ± 1.3%
	3.2%	6.0%	3.8%	1.0%	3.4%
Min-K% Prob	73.8%	76.4%	71.5%	66.8%	64.3%
	70.0 ± 1.5%	72.8 ± 2.3%	68.9 ± 2.2%	64.8 ± 1.4%	59.2 ± 2.4%
	3.8%	3.6%	2.6%	2.0%	5.1%

Can Watermarking Large Language Models Prevent Copyrighted Text Generation and Hide Training Data?

Table 9. AUC of each MIA for the unwatermarked (*top* of each cell), watermarked models (*middle* of each cell) and the drop between the two (*bottom* of each cell) on WikiMIA-64 using UMD scheme.

	Llama-30B	NeoX-20B	Llama-13B	Pythia-2.8B	OPT-2.7B
PPL	66.1%	66.6%	63.6%	58.4%	55.1%
	60.7 ± 3.4%	60.1 ± 3.2%	58.0 ± 3.7%	58.7 ± 1.7%	52.2 ± 2.1%
	5.4%	6.5%	5.6%	-0.3%	2.9%
Lowercase	61.8%	66.4%	62.0%	57.7%	56.6%
	54.8 ± 1.7%	56.8 ± 3.8%	53.8 ± 1.1%	54.5 ± 1.0%	51.4 ± 3.1%
	7.0%	9.6%	8.2%	3.2%	5.2%
Zlib	67.4%	68.1%	65.3%	60.5%	57.7%
	62.4 ± 3.3%	62.0 ± 2.6%	59.9 ± 3.6%	60.9 ± 1.8%	55.5 ± 1.5%
	5.0%	6.1%	4.9%	5.4%	2.2%
Min-K% Prob	68.4%	72.8%	65.9%	61.2%	58.0%
	64.4 ± 2.9%	67.7 ± 3.3%	62.8 ± 3.4%	59.8 ± 0.7%	55.3 ± 2.3%
	4.0%	5.1%	3.1%	1.4%	2.7%

Table 10. AUC of each MIA for the unwatermarked (*top* of each cell), watermarked models (*middle* of each cell) and the drop between the two (*bottom* of each cell) on WikiMIA-32 using UMD scheme.

	Llama-30B	NeoX-20B	Llama-13B	Pythia-2.8B	OPT-2.7B
PPL	69.4%	69.0%	67.5%	61.3%	58.2%
	63.6 ± 5.2%	62.7 ± 3.5%	61.4 ± 5.7%	60.8 ± 2.3%	55.2 ± 2.1%
	5.5%	6.3%	6.1%	0.5%	3.0%
Lowercase	64.1%	68.2%	63.9%	60.9%	59.2%
	54.9 ± 1.8%	59.4 ± 4.8%	54.2 ± 1.8%	55.5 ± 1.6%	52.1 ± 3.9%
	9.2%	8.8%	9.7%	0.6%	2.8%
Zlib	69.8%	69.2%	67.8%	62.1%	59.4%
	64.4 ± 4.7%	63.2 ± 2.8%	62.3 ± 5.1%	61.5 ± 1.9%	56.6 ± 1.6%
	5.4%	6.0%	5.5%	0.6%	2.8%
Min-K% Prob	70.1%	72.1%	67.9%	61.8%	59.2%
	66.2 ± 4.2%	67.1 ± 4.2%	64.5 ± 4.1%	61.0 ± 1.5%	55.8 ± 2.3%
	3.9%	5.0%	3.4%	0.8%	3.4%

Table 11. Results for Smaller Ref attack on WikiMIA-256. The first two rows represent the pair of target and smaller reference model, “No model w.” row represents the baseline AUC of a unwatermarked target LLM and unwatermarked reference model, the other three “double rows” correspond to different variations of the reference model and each cell contains the AUC followed by the AUC drop in comparison to the baseline.

	Llama-30B	NeoX-20B	Llama-13B	Pythia-2.8B	OPT-2.7B
	Llama-7B	Neo-125M	Llama-7B	Pythia-70M	OPT-350M
No model w.	74.7%	70.2%	70.5%	63.6%	64.4%
Ref. not w.	69.7 ± 3.3%	61.0 ± 1.8%	66.3 ± 4.6%	61.6 ± 2.0%	53.2 ± 3.4%
	5.0%	9.2%	4.2%	2.0%	11.2%
Ref. diff. seed	61.7 ± 4.4%	55.5 ± 3.4%	54.1 ± 4.4%	58.3 ± 2.4%	51.3 ± 4.3%
	13.0%	15.0%	16.4%	5.3%	13.1%
Ref. diff. str.	73.7 ± 2.6%	61.0 ± 3.2%	68.8 ± 4.8%	62.5 ± 1.2%	57.3 ± 3.6%
	1.0%	9.2%	1.7%	1.1%	7.1%




Figure 4. AUC drop due to watermarking for each MIA when varying the strength of the watermark.

Table 12. We show the AUC drop when we vary the percentage of green tokens between 30% and 70%. We bold the scenarios when a specific percentage value induces AUC drops for all the attacks.

		0.3	0.4	0.5	0.6	0.7
Llama-30B	PPL	0.71	-0.10	1.40	2.14	1.65
	Lowercase	0.27	2.60	4.24	2.45	3.15
	Zlib	0.58	-0.18	0.66	1.10	0.52
	Min-K% Prob	1.38	0.03	1.28	1.64	1.45
NeoX-20B	PPL	5.56	5.65	6.64	6.92	4.96
	Lowercase	9.84	11.12	12.79	11.94	9.78
	Zlib	6.68	5.94	6.54	6.51	4.71
	Min-K% Prob	0.45	0.88	1.83	1.26	3.84
Llama-13B	PPL	0.19	-0.86	1.22	1.84	1.39
	Lowercase	0.49	2.45	3.93	1.54	1.65
	Zlib	1.31	0.28	1.51	2.00	1.29
	Min-K% Prob	2.81	1.21	2.45	2.70	2.78
Pythia-2.8B	PPL	4.65	4.49	3.39	4.42	3.66
	Lowercase	4.91	6.23	4.18	5.33	7.22
	Zlib	5.37	3.97	3.07	3.17	2.20
	Min-K% Prob	1.10	1.21	1.44	3.42	4.71
OPT-2.7B	PPL	5.45	5.39	5.55	5.18	5.76
	Lowercase	7.71	9.91	9.23	9.83	7.28
	Zlib	3.35	4.12	4.57	4.17	4.11
	Min-K% Prob	2.30	2.07	2.40	3.90	5.00

A.4. Theoretical analysis

Notations and assumptions. We assume that the set of all copyrighted texts C_D that were part of the training data has m elements $\{s_1, s_2, \dots, s_m\}$. Also, we assume that each copyrighted text has a fixed length n , and they are independent from each other.

Theorem A.1. For an LLM watermarked using a “hard” UMD scheme with a percentage of γ green tokens, if $m \cdot \gamma^n < 1$, then the probability of generating at least one copyrighted text from T trials is lower than $m \cdot T \cdot \gamma^n$.

Proof. Given one sample $s = t_1 \oplus t_2 \oplus \dots \oplus t_n \in C_D$. For a “hard” watermarking scheme, the probability $P(s)$ of generating s is smaller than the probability of each token t_i to be on a green list. So, $P(s) < \gamma^n$. The probability of not generating any $s_j \in C_D$ is $P(\neg s_1 \wedge \neg s_2 \wedge \dots \wedge \neg s_m) = \prod_{i=1, m} (1 - P(s_i)) > (1 - \gamma^n)^m > 1 - m\gamma^n$. Note that we used Bernoulli’s inequality at the end. The probability of not generating any $s_j \in C_D$ from T trials is $> (1 - m\gamma^n)^T$, so it is larger than $1 - m \cdot T \cdot \gamma^n$, again, by applying Bernoulli’s inequality. So, the probability of generating at least one copyrighted text from T trials is lower than $1 - (1 - m \cdot T \cdot \gamma^n)$ and hence lower than $m \cdot T \cdot \gamma^n$.

Example. Let’s consider a “hard” UMD watermarking scheme with $\gamma = 0.5$. Let’s assume each copyrighted text is 100 tokens, the model was trained on a dataset containing 10^9 copyrighted texts and we make 10^9 trials to generate copyrighted texts using the LLM. The probability to generate at least one copyrighted text out of these 10^9 trials is $< 10^9 \cdot 10^9 \cdot 0.5^{100} = \frac{10^{18}}{2^{100}} = \frac{1000^6}{1024^{10}} < \frac{1000^6}{1000^{10}} = 1000^{-4} = 10^{-12}$ and hence very low.

Theorem A.2. Let f be a LLM and f_W its watermarked version with a “soft” UMD scheme and let $\epsilon \in (0, \frac{1}{4})$. Let $s = t_1 \oplus t_2 \oplus \dots \oplus t_n \in C_D$ be a copyrighted sample. We consider $\gamma = 0.5$. We denote the output of the softmax layer of f for generating the token t_i as $\frac{a_i}{d_i + a_i}$ and in the case of f_W , we denote it by $\frac{a_i \cdot e^\delta}{b'_i + c'_i \cdot e^\delta + a_i \cdot e^\delta}$ (if t_i is on the green list) and $\frac{a_i}{b''_i + c''_i \cdot e^\delta + a_i}$ (if t_i is on the red list), where a_i is the exponential of the logit value corresponding to the token t_i and b'_i, b''_i and c'_i, c''_i are the sum of the exponentials of the logits corresponding to other tokens that are on the red list and green list, respectively. We assume that $\frac{x}{a_i} < M = \frac{1-4\epsilon}{1+4\epsilon}$, for all $x \in \{d_i, b'_i, b''_i, c'_i, c''_i\}$ which would restrict f to be relatively confident in its predictions for each token t_i . Then, we can always find a δ (strength) for the watermarking scheme such that the probability of generating s is reduced by at least $(1 + \frac{2\epsilon}{2\epsilon+1})^n$ times in comparison to the case of the unwatermarked model.

Proof. First, we observe that the probability of generating the token t_i by the unwatermarked model is $\frac{a_i}{d_i + a_i} = \frac{1}{\frac{d_i}{a_i} + 1} > \frac{1}{M+1} = 1/2 + 2\epsilon$.

We observe that since there is a finite number of $\frac{x}{a_i}$ and they are all positive, then it exist a lower bound for $\frac{x}{a_i}$ (let’s denote it by $m > 0$). Since $\gamma = 0.5$, the probability of t_i being a green token is $\frac{1}{2}$ and hence the probability of the watermarked model to generate t_i is $\frac{1}{2} \frac{a_i \cdot e^\delta}{b'_i + c'_i \cdot e^\delta + a_i \cdot e^\delta} + \frac{1}{2} \frac{a_i}{b''_i + c''_i \cdot e^\delta + a_i} < \frac{1}{2} + \frac{1}{2} \frac{a_i}{b''_i + c''_i \cdot e^\delta + a_i} = \frac{1}{2} + \frac{1}{2} \frac{1}{\frac{b''_i}{a_i} + \frac{c''_i}{a_i} \cdot e^\delta + 1} \leq \frac{1}{2} + \frac{1}{2} \frac{1}{m \cdot (e^\delta + 1) + 1}$. We pick $\delta > \log(\frac{1-2\epsilon(m+1)}{2\epsilon m})$ and we observe that $\frac{1}{2} + \frac{1}{2} \frac{1}{m \cdot (e^\delta + 1) + 1} < \frac{1}{2} + \frac{1}{2} \frac{1}{m \cdot (\frac{1-2\epsilon(m+1)}{2\epsilon m} + 1) + 1} = \frac{1}{2} + \frac{1}{2} \frac{1}{m \cdot (\frac{1-2\epsilon}{2\epsilon m}) + 1} = \frac{1}{2} + \epsilon$.

So, by combining the two observations above, we conclude that the probability of generating t_i is reduced by at least $\frac{\frac{1}{2} + 2\epsilon}{\frac{1}{2} + \epsilon} = 1 + \frac{2\epsilon}{2\epsilon+1}$ times. Therefore, since there are n tokens in s , the probability of generating s is reduced by at least $(1 + \frac{2\epsilon}{2\epsilon+1})^n$ times.

Observation. Since the probability is reduced by at least $(1 + \frac{2\epsilon}{2\epsilon+1})^n$ times in Theorem A.2 then the probability of generating s is lower than $(\frac{2\epsilon+1}{4\epsilon+1})^n$ (as the maximum probability of generating with the unwatermarked model is 1). Hence, with the notations from Theorem A.1, the probability of generating at least one copyrighted text from T trials is lower than $m \cdot T \cdot (\frac{2\epsilon+1}{4\epsilon+1})^n$

Takeaways. Our theoretical analysis demonstrates that watermarking significantly reduces the probability of generating copyrighted text verbatim. For both a “hard” and “soft” UMD scheme, the upper bound for the likelihood of producing copyrighted content from a set of trials decreases exponentially with the length of the copyrighted texts.

A.5. Improving Detection Performance with Adaptive Min-K% Prob

This section demonstrates how an informed, adaptive attacker can improve the success rate of a recent MIA, Min-K% Prob. Our main idea is that if the attacker has some knowledge about the watermarking technique (green-red token lists and watermark’s strength δ), then they can readjust the token probabilities even if they don’t have any additional information about the logit distribution other than the probability of each token from the target sample given the previous ones. Note that having this knowledge about the watermarking scheme is an assumption made by prior work for public watermark detection purposes (Kirchenbauer et al., 2023) while having access to the probability of each token from a sample given the previous ones is the assumption Min-K% Prob method makes as well.

Our method described in Algorithm 1 is based on the observation that if the denominator of softmax function (i.e., $\sum_i e^{z_i}$, where z_i is the logits for the i -th vocabulary) does not vary significantly between watermarked samples and unwatermarked samples, then we can readjust the probabilities of the green tokens by “removing” the bias δ . More precisely, assuming the approximation for the denominator of softmax is good, then the probability for each token t_i in an unwatermarked model will be around $\frac{e^{L_i}}{c}$, where L_i is the logit corresponding to the token t_i and c is a constant. However, for a watermarked model, if the token t_i is green, then the probability would be approximated by $\frac{e^{L_i + \delta}}{d}$, where d is again a constant, while in the case t_i is red the probability will be around $\frac{e^{L_i}}{d}$. To compensate for the bias introduced by watermarking, we divide

Algorithm 1 Adaptive Min-K% Prob

Require : Tokenized target sample $t = t_1 \oplus t_2 \oplus \dots \oplus t_n$, access to the probability of the target (watermarked) LLM f to generate t_i given the $i - 1$ previous tokens and t_0 (empty string) $f(t_i|t_0 \oplus t_1 \oplus \dots \oplus t_{i-1})$ (similar assumption as Min-K% Prob algorithm), K , we assume we know the watermarking scheme (e.g., for public watermark detection purposes), i.e. we know the green and red lists as well as δ .

Output : Adjusted average of the minimum $K\%$ token probabilities when generating $t_1 \oplus t_2 \oplus \dots \oplus t_n$

```

adj_prob ← {}
for  $i \in 1, 2, \dots, n$  do
     $p_f(t_i) \leftarrow f(t_i|t_0 \oplus t_1 \oplus \dots \oplus t_{i-1})$ 
    if  $t_i$  is green then
        adj_prob ← adj_prob  $\cup$  {  $\frac{p_f(t_i)}{e^\delta}$  }
    else
        adj_prob ← adj_prob  $\cup$  {  $p_f(t_i)$  }
    end
end
 $k = \text{floor}(n \cdot K\%)$ 
adj_k_prob ←  $\text{min}_k(\text{adj\_prob})$ 
return mean(log(adj_k_prob))

```

▷ The set of adjusted probabilities
 ▷ Find the number of token probabilities to keep
 ▷ Select the minimum k probabilities
 ▷ Return the mean of the minimum k log-probabilities

Table 13. We show the AUC of Min-%K Prob (referred as “Not adapt.”) and our method (referred as “Adapt.”) when using UMD watermarking scheme. We highlight the cases when our method improves over the baseline.

		Llama-30B	NeoX-20B	Llama-13B	Pythia-2.8B	OPT-2.7B
WikiMIA 32	Not adapt.	66.2%	67.1%	64.5%	61.0%	55.7%
	Adapt.	68.5%	71.3%	66.3%	61.0%	59.1%
WikiMIA 64	Not adapt..	64.4%	67.7%	62.8%	59.8%	55.3%
	Adapt.	67.3%	72.0%	64.9%	60.6%	57.4%
WikiMIA 128	Not adapt..	70.0%	73.0%	68.9%	64.8%	59.2%
	Adapt.	73.1%	75.9%	71.0%	66.4%	64.0%
WikiMIA 256	Not adapt..	70.5%	76.2%	70.4%	69.5%	63.1%
	Adapt.	71.3%	78.2%	72.4%	70.7%	66.2%

the probability of green tokens by e^δ and this way we end up with probabilities that are just a scaled (by $\frac{e}{d}$) version of the probabilities from the unwatermarked model. The scaling factor will not affect the orders between the samples when computing the average of the minimum $K\%$ log-probabilities as long as the tested sentences are approximately the same length, which is an assumption made by Shi et al. (2023) as well.

Despite the strong assumption we assumed regarding the approximation of the denominator, empirical results show that our method effectively improves the success rate of Min-K% under watermarking. The results in Table 13, averaged over 5 runs, demonstrate that our method improves over the baseline in 95% of the cases, and the increase is as high as 4.8%.

Takeaways. We demonstrate that an adaptive attacker can leverage the knowledge of a watermarking scheme to increase the success rate of a recent MIA, Min-K% Prob.

A.6. Additional results on BookMIA

Table 14. Measuring the reduction in verbatim memorization of training texts on BookMIA. We report the relative increase in perplexity between the watermarked and unwatermarked models, where larger values correspond to less memorization. Note that “P.” stands for “prompt length”.

		Llama-30B		Llama-13B	
		P.	Min. Avg.	Min.	Avg.
UMD	0	1.5	33.7	2.4	41.2
	10	1.5	33.6	2.3	41.0
	20	1.4	33.5	2.3	40.8
	100	1.3	32.9	1.9	40.3
Unigram	0	1.6	36.4	2.4	44.5
	10	1.6	36.3	2.4	44.3
	20	1.5	36.1	2.3	44.2
	100	1.4	35.5	1.8	43.6

Table 15. AUC of each MIA for the unwatermarked (*top* of each cell), watermarked models (*middle* of each cell), and the drop between the two (*bottom* of each cell) on BookMIA using UMD scheme.

	Llama-30B	Llama-13B
PPL	85.4%	68.2%
	84.7 ± 1.4%	67.6 ± 2.5%
	0.7%	0.6%
Lowercase	87.9%	77.6%
	80.9 ± 3.1%	67.2 ± 4.0%
	7.0%	10.4%
Zlib	82.5%	62.5%
	77.8 ± 1.2%	57.1 ± 2.0%
	4.7%	5.4%
Min-K% Prob	85.1%	70.2%
	85.0 ± 1.0%	68.5 ± 0.1%
	0.1%	1.7%