

POSTMARK: A Robust Blackbox Watermark for Large Language Models

Yapei Chang[✉] Kalpesh Krishna[✉]
Amir Houmansadr[✉] John Wieting[✉] Mohit Iyyer[✉]

[✉]University of Massachusetts Amherst, [✉]Google
{yapeichang, amir, miyyer}@cs.umass.edu
{kalpeshk, jwieting}@google.com

Abstract

The most effective techniques to detect LLM-generated text rely on inserting a detectable signature—or *watermark*—during the model’s decoding process. Most existing watermarking methods require access to the underlying LLM’s logits, which LLM API providers are loath to share due to fears of model distillation. As such, these watermarks must be implemented independently by each LLM provider. In this paper, we develop POSTMARK, a modular post-hoc watermarking procedure in which an input-dependent set of words (determined via a semantic embedding) is inserted into the text *after* the decoding process has completed. Critically, POSTMARK does not require logit access, which means it can be implemented by a third party. We also show that POSTMARK is more robust to paraphrasing attacks than existing watermarking methods: our experiments cover eight baseline algorithms, five base LLMs, and three datasets. Finally, we evaluate the impact of POSTMARK on text quality using both automated and human assessments, highlighting the trade-off between quality and robustness to paraphrasing. We release our code, outputs, and annotations at <https://github.com/lilakk/PostMark>.

1 Introduction

Large language models (LLMs) are increasingly being deployed for malicious applications such as fake content generation. The consequences of such applications for the web as a whole are dire: modern LLMs are known to hallucinate (Xu et al., 2024), and their outputs may contain biases and artifacts that are a product of their training data (Navigli et al., 2023). If the web is flooded with millions of LLM-generated articles, how can we trust the veracity of the content we are reading? Additionally, do we want to train LLMs of the future on text generated by LLMs of the present (Shumailov et al., 2023)?

To combat this emerging problem, researchers have developed several *LLM-generated text detection* techniques that leverage watermarking (Aaronson and Kirchner, 2022; Kirchenbauer et al., 2023), outlier detection (Mitchell et al., 2023), trained classifiers (Tian, 2023), or retrieval-based methods (Krishna et al., 2023). Among these, watermarking methods that embed detectable signatures into model outputs tend to be the most effective and robust (Krishna et al., 2023). However, most watermarking algorithms require access to the logits of the underlying LLM, which means that they can only be implemented by individual LLM API providers such as OpenAI or Google (Yang et al., 2023). Furthermore, while these methods are able to achieve high detection rates with minimal false positives, their effectiveness goes down when the LLM-generated text is modified through paraphrasing, translation, or cropping (Krishna et al., 2023; He et al., 2024; Kirchenbauer et al., 2024).

In this work, we develop POSTMARK, a watermarking approach with relatively high detection rates even in the presence of paraphrasing attacks. POSTMARK is a **post-hoc watermark** that given some model-generated text, finds words conditioned on the semantics of the text using an embedding model, then calls a separate instruction-following LLM to insert these words into the text without appreciably modifying its meaning. Unlike prior methods, POSTMARK only requires access to the outputs of the underlying LLM (i.e., no logits).

Overall, our contributions are threefold: **1.** We propose POSTMARK, a novel post-hoc watermarking method that can be applied by third-party entities to outputs from an API provider like OpenAI. **2.** We conduct extensive experiments across eight baseline algorithms, five base LLMs, and three datasets, showing that POSTMARK offers superior robustness to paraphrasing attacks compared to existing methods. **3.** We verify through a human evaluation that the words inserted by POSTMARK

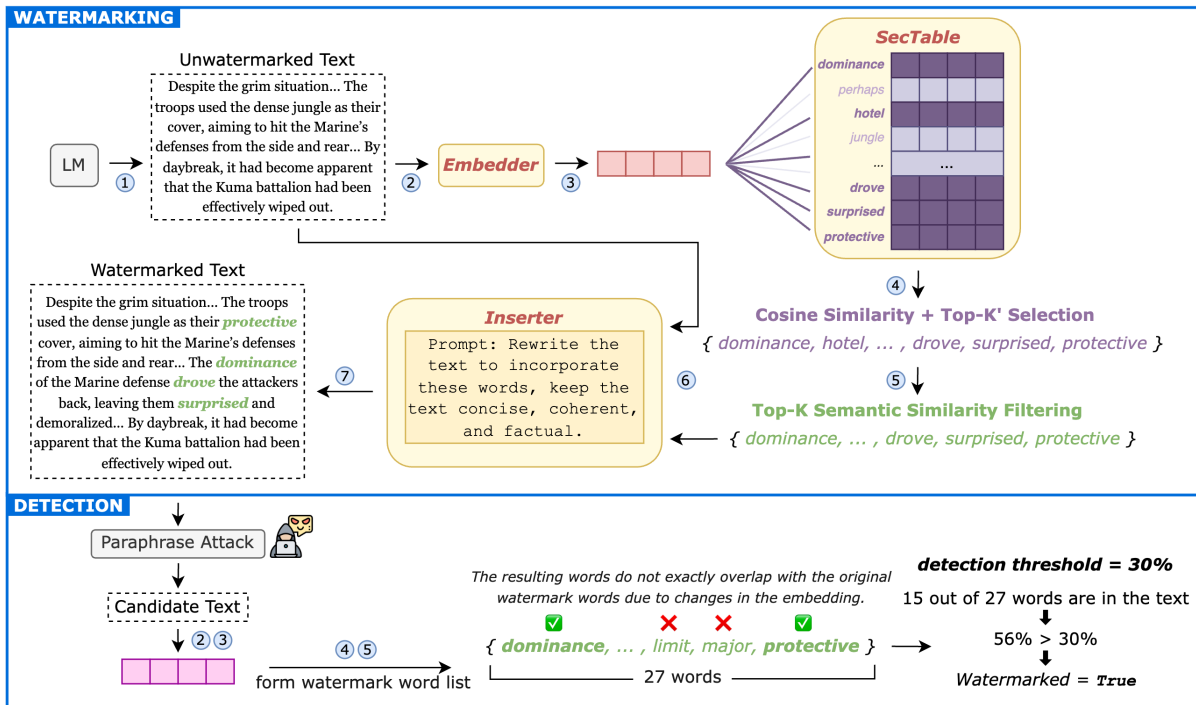


Figure 1: The POSTMARK watermarking and detection procedure. Given some unwatermarked input text, we generate its embedding using the EMBEDDER and compute its cosine similarity with all word embeddings in the SECTABLE, performing top-k selection and additional semantic similarity filtering to choose a list of words. Then, we instruct the INSERTER to watermark the text by rewriting it to incorporate all selected words. During detection, we similarly obtain a watermark word list and check how many of these words are present in the input text.

during watermarking cannot be reliably detected by humans. We also conduct comprehensive quality evaluations encompassing *coherence*, *relevance*, and *interestingness* for various watermarking methods. Notably, we also assess *factuality*, an aspect that has not been evaluated in prior work. Our findings reveal that relatively robust watermarks all negatively affect factuality.

2 POSTMARK: a post-hoc watermark

Most existing watermarking algorithms embed the watermark during the LLM’s decoding process. For example, the watermark of Kirchenbauer et al. (2023, KGW) partitions an LLM’s vocabulary into two lists (a *green* list and a *red* list) at each decoding timestep based on a hash of the previous word, and then upweights the green tokens such that they are more likely to be sampled than red tokens. These watermarks have several issues: (1) they require access to the LLM’s logits; (2) because they rely on modifications to the next-token probability distribution, their effectiveness diminishes on LLMs that produce lower-entropy distributions, such as those that have undergone RLHF (Bai et al., 2022); and (3) they show limited robustness to paraphrasing attacks as demonstrated by our results

in Section 3.2 and supported by findings from prior work (Krishna et al., 2023; Sadasivan et al., 2024).

In response, we develop POSTMARK, a watermarking method that does not require logit access, maintains high detection rates on low-entropy models and tasks, and exhibits improved robustness to paraphrasing attacks. Unlike existing watermarks, POSTMARK requires access to just the text generated by the underlying LLM, not the next-token distributions. The rest of this section fully specifies POSTMARK’s operation.

Intuition and terminology: At a high level, POSTMARK is based on the intuition that a text’s semantics should not drastically change after watermarking or paraphrasing. Thus, we can condition our watermark on a semantic embedding of the input text that ideally changes only minimally when paraphrasing is applied. To make this work, we rely on three modules: an embedding model EMBEDDER, a secret word embedding table SECTABLE, and an insertion model INSERTER implemented via an instruction-following LLM.

Figure 1 illustrates POSTMARK’s watermarking and detection pipelines. First, we generate the embedding of an input text using the EMBEDDER. We then compute the cosine similarity between

this embedding and all of the word embeddings in SECTABLE, performing top- k selection and filtering to form a watermark word list. Next, we prompt INSERTER to smoothly incorporate the selected words into the input to create the watermarked text. During detection, we follow similar steps to obtain a word list, and check how many of the words are present in the input text.

Embedding model EMBEDDER: The EMBEDDER needs to be capable of projecting both words and documents into a high-dimensional latent space. In our main experiments, we use OpenAI’s TEXT-EMBEDDING-3-LARGE (OpenAI, 2024b), a powerful model that demonstrates strong performance on the MTEB benchmark (Muennighoff et al., 2023). However, any embedding model can be used here. In Section 3.2, we also experiment with NOMIC-EMBED (Nussbaum et al., 2024), an open-source model.

Secret word embedding table SECTABLE: The core idea behind POSTMARK is to use an LLM to insert a list of watermark words into the input text without appreciably modifying the quality or meaning of the text, where the words in the list are selected by computing the cosine similarity between the text embedding and a word embedding table SECTABLE. The construction of SECTABLE involves two main steps, which we detail below:

> **Step 1. Choosing a vocabulary \mathbb{V} :** To decide which words to include in SECTABLE, we use the WikiText-103 corpus (Merity et al., 2017) as our base vocabulary. To avoid inserting arbitrary words that make little sense, we remove all function words, proper nouns, and infrequent rare words. This refined set forms our final vocabulary, \mathbb{V} . We provide more details on this filtering process in §A. > **Step 2. Mapping words in \mathbb{V} to embeddings:** To make it difficult for attackers to recover our embedding table, we construct SECTABLE by randomly assigning each word in the vocabulary to an embedding produced by EMBEDDER; the resulting mapping acts as a cryptographic key.¹ More specifically, we generate a set of embeddings \mathbb{D} for a collection of random documents using EMBEDDER and then randomly map each word in \mathbb{V} to a unique document embedding in \mathbb{D} to produce SECTABLE.²

¹We could also just use EMBEDDER’s word embeddings as SECTABLE directly. However, this can easily be recovered by an attacker, and our experiments show that it also reduces POSTMARK’s effectiveness due to many words already being present in the input text even before insertion.

²The selection of these documents is flexible. In our ex-

Insertion model INSERTER: The INSERTER needs to have instruction-following capabilities, and its purpose is to rewrite the input text to incorporate words from the watermark word list. We use GPT-4o (OpenAI) as the INSERTER in our main experiments, and later show in Section 3.2 that open-source models like LLAMA-3-70B-INST (AI@Meta, 2024) also show promising performance.

2.1 Inserting the watermark

> **Step 1. Deciding how many words to insert:** How many words should we insert into a given text? We define a hyperparameter called the **insertion ratio** r that determines this number. The insertion ratio represents the percentage of the input text’s word count: for example, if $r = 10\%$ and the input text has 50 words, we will insert 5 words.

> **Step 2. Obtaining a watermark word list:** Suppose that the watermark list should contain k words. To create the watermark word list given the input text, we first compute the input’s embedding $e_t = \text{EMBEDDER}(t)$. Next, we compute $\text{CosineSimilarity}(e_t, \text{SECTABLE})$ and select the top k' most similar words, then perform semantic similarity filtering to obtain the final k words.³ We present an analysis on how frequently a word is chosen as a watermark word in §A.

> **Step 3. Inserting words into the text:** To watermark the text, we instruct INSERTER to rewrite it via zero-shot prompting, incorporating words in the watermark word list while keeping the rewritten text coherent, factual, and concise.⁴ The prompt can be found in §B.

2.2 Detecting the watermark

During detection, given some text, the goal is to find out if the text contains a watermark. Similar to the watermarking procedure, we embed the candidate text using EMBEDDER, form a word list, and then check how many words in the list are present in the text by computing a presence score p :

periments, we randomly sample 250-word snippets from the RedPajama dataset’s English split (Computer, 2023).

³Due to the random nature of the word-to-embedding mapping of T , the top k' words might include highly irrelevant words (e.g., “hotel” in Figure 1). Thus, we refine the top- k' list by selecting the top k words whose *actual* embeddings (as determined by EMBEDDER) are most similar to e_t .

⁴In practice, we find that dividing a long word list into sublists of 10 words each and then iteratively asking the INSERTER to incorporate each sublist ensures a high insertion success rate. This may not be necessary if the INSERTER has better instruction-following capabilities.

$$p = \frac{|\{w \in \text{list s.t. } \exists w' \in \text{text}, \text{sim}(w', w) \geq 0.7\}|}{|\text{list}|}$$

A word w is marked present in the text if there is any other word w' with an embedding cosine similarity greater than a threshold that we set to 0.7. We choose this method over exact match to ensure additional robustness against paraphrasing.⁵ If p is larger than a certain threshold, it is likely that the text has been watermarked. As later discussed in Section 3.1, the primary metric we use to measure detection accuracy is the true positive rate at a fixed 1% false positive rate. We thus set the threshold to ensure a 1% FPR, same as what we do for all baselines in our main experiments.

3 Experiments

In this section, through extensive experiments on three datasets and five language models, we demonstrate that POSTMARK consistently outperforms both logit-free and logit-based methods in terms of robustness to paraphrasing attacks, especially on low-entropy models that have undergone RLHF alignment. Furthermore, we showcase POSTMARK’s modular design by testing an open-source variant, which achieves promising results.

3.1 Experimental setup

Baselines: We compare POSTMARK against 8 baseline algorithms, more detailed descriptions can be found in §D. (1) **KGW** (Kirchenbauer et al., 2023): Partitions the vocabulary into “green” and “red” lists based on the previous token, then boosts the probability of green tokens during generation. (2) **Unigram** (Zhao et al., 2023): A more robust variant of KGW that uses a fixed partition for all generations. (3) **EXP** (Aaronson and Kirchner, 2022): Uses exponential sampling to bias token selection with a pseudo-random sequence. (4) **EXP-Edit** (Kuditipudi et al., 2024): A variant of EXP that uses edit distance during detection. (5) **Sem-Stamp** (Hou et al., 2023): A sentence-level algorithm that partitions the sentence semantic space. (6) **k-SemStamp** (Hou et al., 2024): Improves SemStamp by using k-means clustering to partition the

semantic space. (7) **SIR** (Liu et al., 2024b): Generates watermark logits from the semantic embeddings of preceding tokens then adds them to the model’s logits. (8) **Blackbox** (Yang et al., 2023): This method, like ours, works in a blackbox setting where only model outputs are visible. It substitutes words representing bit-0 in a binary encoding scheme with synonyms representing bit-1.

Hyperparameters: The key hyperparameter for POSTMARK is the insertion ratio r , which controls how many words are inserted during the watermarking process. We set r to 12% as preliminary experiments suggest that this value strikes a good balance between quality and robustness. Section 4.1 explores different POSTMARK configurations that vary r . In all following discussion and tables, we refer to these configurations with the naming convention “POSTMARK@ r ”. We carefully tune all baselines’ hyperparameters to maximize their robustness to paraphrasing; more details in §D.

Base models: Our experiments involve five generative models: LLAMA-3-8B (AI@Meta, 2024), LLAMA-3-8B-INST (AI@Meta, 2024), MISTRAL-7B-INST (Jiang et al., 2023), GPT-4 (OpenAI, 2024a), and OPT-1.3B (Zhang et al., 2022). Among these, LLAMA-3-8B-INST, MISTRAL-7B-INST, and GPT-4 have been aligned with human preferences. For details on model checkpoints and generation length, see §E. We do not run OPT-1.3B ourselves but directly use its unwatermarked outputs provided by Hou et al. (2024). Due to difficulties in running SemStamp, k-SemStamp, and SIR,⁶ we apply POSTMARK directly to these outputs and compare our results with the published numbers in Hou et al. (2024).

Datasets: Our main experiments use three datasets: (1) **OpenGen**, a dataset collected by Krishna et al. (2023) designed for open-ended generation that consists of two-sentence chunks sampled from the validation set of WikiText-103; (2) **LFQA**, a dataset collected by Krishna et al. (2023) for long-form question answering that contains questions sampled from the *r/explainlikeimfive* subreddit that span multiple domains; and (3) **RealNews** (Raffel et al., 2020), a subset of the C4 dataset that includes news articles gathered from a wide range of reliable news websites.

⁵We use the PARAGRAM word embedding model developed by Wieting et al. (2015) to compute cosine similarity for this step. This model is chosen for its superior performance in assigning high similarity scores to close synonyms and low scores to unrelated words, more details in §C.

⁶Their code is available but not runnable yet. We look forward to running these methods ourselves once the issues are resolved.

Paraphrasing attack setup: Following prior work (Hou et al., 2023, 2024; Kirchenbauer et al., 2024; Liu et al., 2024b), we use GPT-3.5-TURBO as our paraphraser. We use a sentence-level paraphrasing approach where the model iterates through each sentence of the input text, using all preceding context to paraphrase the current sentence. See §F for more details on this setup.

Metric for measuring detection performance: In addition to the true positive rate, a low false positive rate is critical for LLM-generated detection. Thus, following prior detection work (Krishna et al., 2023; Zhao et al., 2023; Hou et al., 2023, 2024; Liu et al., 2024b), we use **TPR at 1% FPR** as our primary metric.

3.2 Results

We present our main experimental results on robustness to paraphrasing attacks in Table 1, and discuss our main findings below. Runtime analysis and API cost estimates can be found in §G.

POSTMARK is an effective and robust watermark. POSTMARK consistently achieves a high TPR *before paraphrasing* (> 90%), outperforming baselines like Blackbox, KGW, and EXP-Edit. Additionally, POSTMARK achieves higher TPR *after paraphrasing* compared to other baselines, including Blackbox, the only other method that operates under the same logit-free condition. The only settings that POSTMARK is not the most robust model under paraphrasing is with LLAMA-3-8B and OPT-1.3B, where Unigram and EXP respectively exhibit more robustness. We note that Unigram is much more vulnerable to reverse-engineering than POSTMARK because it uses a fixed *green/red* list partition for all inputs, which can be exploited with repetition attacks.⁷ Both Unigram and EXP’s effectiveness diminish with low-entropy models. In §I and §J, we also observe that both methods significantly degrade text quality, leading to excessive repetitions.

Logit-based baselines perform worse on low-entropy models and tasks, while POSTMARK stays relatively unaffected. Results from Table 1 demonstrate that logit-based baselines (i.e.,

⁷For Unigram, detection works by comparing the number of *green* tokens present in the input text to the expected count under the null hypothesis of no watermarking. The adversary can pick a word “apple” and submit a long repeating sequence of this word (e.g., “apple apple apple...”) to the watermark detection service. If it says this sequence is watermarked, then “apple” must be in the green list.

all baselines except Blackbox) generally perform worse on aligned models (LLAMA-3-8B-INST and MISTRAL-7B-INST) compared to the non-aligned LLAMA-3-8B, and worse on LFQA than on Open-Gen. This performance difference is consistent with findings from prior work (Kuditipudi et al., 2024) and can be attributed to the lower entropy of aligned models resulting from RLHF or instruction-tuning, as well as the inherently lower entropy of the LFQA task. The “Avg Entropy” column of Table 1 illustrates these entropy differences. In contrast, POSTMARK consistently outperforms all baselines in terms of robustness against paraphrasing attacks in these low-entropy scenarios.

Open-weight POSTMARK shows promise. While our main experiments use GPT-4O as the INSERTER and OpenAI’s TEXT-EMBEDDING-3-LARGE as the EMBEDDER, we show in Table 2 that an open-weight combination of LLAMA-3-70B-INST and NOMIC-EMBED can also achieve promising robustness to paraphrasing attacks. The modular design of POSTMARK allows for flexible experimentation with various components. As each module’s capabilities advance, POSTMARK’s robustness will likewise improve.

4 Impact of watermarking on text quality

POSTMARK modifies text during watermarking by inserting new words, which often results in longer watermarked text.⁸ Table 3 shows several common types of edits made by POSTMARK during watermarking.⁹ Although edits adding new content are expected to hurt quality, this quality degradation is not unique to POSTMARK. Prior work has found that all watermarking methods negatively affect text quality to some extent (Singh and Zou, 2023). For logit-based methods like KGW, quality degradation occurs because relevant words can be downweighted during decoding. While existing papers on watermarking often lack extensive quality evaluations, we conduct both automatic and human evaluations to assess the quality of watermarked text (*relevance, coherence, interestingness, and factuality*) in this section.

Semantic meaning preservation To check whether POSTMARK preserves the general semantic meaning of the original unwatermarked text, we compute the average cosine similarity between the

⁸A full table of length comparison is in §H.

⁹Summarized based on a small-scale qualitative analysis.

Metric → Model ↓	Dataset ↓	Avg Entropy ↓	TPR at 1% FPR (Before Paraphrasing / After Paraphrasing)								
			POSTMARK@12	Blackbox	KGW	Unigram	EXP	EXP-Edit	SIR	SemStamp	k-SemStamp
LLAMA-3-8B	OpenGen	3.6	99.7 / 63.5	81.2 / 2.2	100 / 74.8	99.8 / 93.4	99.8 / 36.6	97.3 / 73.3	-	-	-
	LFQA	3.5	97.8 / 72.5	82.8 / 1.6	99.8 / 25.6	99.8 / 79.6	99.8 / 12.4	83 / 41	-	-	-
LLAMA-3-8B-INST	OpenGen	1.6	99.4 / 46.4	91.8 / 1	98.2 / 21.6	99.6 / 41.4	99.6 / 4.8	47.8 / 2.2	-	-	-
	LFQA	1.3	96 / 65.7	86.2 / 3	85.8 / 19	98.6 / 31.8	98.4 / 0.6	21.1 / 0.6	-	-	-
MISTRAL-7B-INST	OpenGen	1.4	99.2 / 69.2	98.4 / 0.4	100 / 16	99.8 / 56	99.4 / 5	33 / 1.5	-	-	-
	LFQA	1.1	99.6 / 56.4	89.8 / 0.4	99.4 / 23.6	97.2 / 41.2	97.4 / 0.8	20.1 / 2.1	-	-	-
GPT-4	OpenGen	-	99.4 / 59.4	99.4 / 1.4	-	-	-	-	-	-	-
	LFQA	-	99.4 / 65	99.2 / 0.4	-	-	-	-	-	-	-
OPT-1.3B	RealNews	3.6	98.2 / 67.2	1.2 / 0	99.2 / 40.8	98.8 / 77.2	99.4 / 80.7	69.6 / 46.7	99.4 / 24.7	93.9 / 33.9	98.1 / 55.5

Table 1: Comparison of POSTMARK and baselines. All numbers are computed over 500 generations. Each entry shows the TPR at 1% FPR **before paraphrasing** and **after paraphrasing**. The “Avg Entropy” column shows the average token-level entropy (in bits) of each model on each dataset.

POSTMARK@12 Impl.	TPR at 1% FPR
Closed	99.4 / 59.4
Open	100 / 52.1

Table 2: TPR at 1% FPR **before** and **after paraphrasing**. The open-source implementation of POSTMARK@12 with NOMIC-EMBED as the EMBEDDER and LLAMA-3-70B-INST as the INSERTER shows promising performance on OpenGen with GPT-4 as the base LLM.

embeddings unwatermarked and watermarked outputs in Table 4, and find the similarity score to be consistently around 0.95.

Setting up quality evaluations: Prior work on watermarking has predominantly used perplexity as a measure for text quality (Kirchenbauer et al., 2023; Zhao et al., 2023; Yang et al., 2023; Liu et al., 2024b; Hu et al., 2024; Hou et al., 2023, 2024). However, perplexity alone has been shown to be an unreliable indicator of quality (Wang et al., 2022). Some studies have explored alternative methods, such as LLM-based evaluations (Singh and Zou, 2023) and human assessments (Kirchenbauer et al., 2024). Here, we evaluate the quality of watermarked text using automated and human evaluations, aiming to address four key questions:

- > **Q1:** How does POSTMARK compare to other baselines in terms of impact on text quality?
- > **Q2:** What is the quality-robustness trade-off for POSTMARK?
- > **Q3:** How often do humans think that POSTMARK watermarked texts are at least as good as their unwatermarked versions?
- > **Q4:** Are words inserted by POSTMARK easily detectable by humans?

4.1 Automatic evaluation

In this section, we compare POSTMARK with other baselines regarding impact on quality (**Q1**) and

address the quality-robustness trade-off of POSTMARK (**Q2**).

Pairwise preference evaluation setup: We adopt the LLM-as-a-judge (Zheng et al., 2023) setup to perform a pairwise comparison task. We choose GPT-4-TURBO as our judge as it is the high-ranked evaluator model on the Reward Bench leaderboard (Lambert et al., 2024)¹⁰ that we can easily access. Given 100 OpenGen prefixes and corresponding pairs of anonymized unwatermarked and watermarked responses, the model evaluates each pair and chooses which response it prefers, where ties are allowed. The model is instructed to consider the *relevance*, *coherence*, and the *interestingness* of the responses when making a judgment. The full prompt can be found in §K. Then, we compute the **soft win rate** of various baselines in Table 6 and several POSTMARK configurations in Table 7, which equals the number of ties plus the number of wins for the watermarked response.

Factuality evaluation setup: To assess *factuality*, an essential aspect not addressed in the previous pairwise comparisons or previous watermarking research, we use FactScore (Min et al., 2023), an automatic metric that measures the percentage of atomic claims in an LLM-generated biography that are supported by Wikipedia. We generate biographies for the entities in the FactScore dataset and comparing the FactScores of the outputs before and after watermarking. Results are reported in Table 5. Before watermarking, LLAMA-3-8B-INST achieves a score of 40.2. Running KGW, Unigram, POSTMARK@12, and POSTMARK@6 all results in slight reductions in FactScore. Overall, less robust methods (KGW and POSTMARK@6) have

¹⁰The current leaderboard is hosted on [huggingface](https://huggingface.com). GPT-4-TURBO’s high ranking indicates that it is a relatively robust and reliable LLM evaluator.

TYPE	BEFORE WATERMARK	AFTER WATERMARK
Rewriting existing content		
Rewording	<i>Her decision to quit the opera, however, did not lessen the engulfing sadness which veiled her once radiant joy.</i>	<i>Her decision to resign from the opera, however, did not lessen the engulfing sadness which veiled her once radiant joy.</i>
Clarification	<i>Since the charges concerned violation of civil rights and not actual murder, the defendants received surprisingly light sentences, ranging from three to ten years.</i>	<i>Since the charges concerned violation of civil rights and not actual murder, the defendants received surprisingly light sentences, ranging from three to ten years of imprisonment.</i>
Adding new content		
Metaphors	<i>In fact, despite Mount Elbert’s somewhat minimal precipitation, it displays a remarkable ability to sustain life.</i>	<i>In fact, despite Mount Elbert’s somewhat minimal precipitation, it displays a remarkable ability to sustain life, almost as if it wears an armor of resilience, immune to the challenges it faces.</i>
Interpretive claims	<i>He swiftly plants timed explosives around the warehouse, ensuring to place a few on the largest weapon caches for maximum effect.</i>	<i>He swiftly plants timed explosives around the warehouse, ensuring to place a few on the largest weapon caches for maximum effect. The depth of his planning was a testament to his expertise in defense tactics.</i>
New details	<i>Headlam had the ability to foster a culture of discipline, camaraderie and respect among the airmen under his command, reflecting his firm belief in focused team effort and mutual support.</i>	<i>Headlam had the ability to foster a culture of discipline, camaraderie and respect among the airmen under his command, reflecting his firm belief in focused team effort and mutual support. His attention to detail was evident in every aspect of the unit’s operations.</i>

Table 3: Example edits made by POSTMARK during the watermarking process. Changes are highlighted in orange, and watermark words are in bold.

Base LLM	SIM	FactScore
LLAMA-3-8B	94.2	LLAMA-3-8B-INST 40.2
LLAMA-3-8B-INST	94.8	+ KGW 37.8
MISTRAL-7B-INST	94.6	+ Unigram 37.2
GPT-4	95.3	+ POSTMARK@12 37.3
		+ POSTMARK@6 38.3

Table 4: Average cosine similarity between the embeddings of unwatermarked and POSTMARK@12 watermarked outputs on OpenGen. Embeddings are obtained using TEXT-EMBEDDING-3-LARGE. Numbers are averaged over 500 pairs.

Table 5: FactScore evaluation results based on 100 generations with LLAMA-3-8B-INST as the base generator LLM. All four evaluated methods impact factuality negatively to some extent, with less robust methods causing a lesser negative impact.

less negative impact on factuality.

> Q1: POSTMARK does not affect quality as much as other baselines. Results from Table 6 show that POSTMARK performs exceptionally well in pairwise comparisons across models. In contrast, despite Unigram’s strong robustness to paraphrasing—sometimes even outperforming POSTMARK when tested on LLAMA-3-8B—it has a significantly lower soft win rates, especially on LLAMA-3-8B (17%). This low score is likely due to frequent repetitions in Unigram outputs, as detailed in §I. Regarding factuality, KGW, Unigram, and POSTMARK@12 all show similar levels of negative impact as their FactScores are respectively 37.8, 37.2, and 37.3.

> Q2: Inserting more words enhances robustness but hurts quality, and vice versa. We first use the pairwise comparison setup to evaluate the quality-robustness trade-off of POSTMARK with r set to six different values: 6, 8, 12, 15, 20, and 30. Results in Table 7 reveal a strong inverse correlation between quality and robustness, with a Pearson coefficient of -0.98. FACTSCORE@6 also achieves a higher FactScore (38.3) than FACTSCORE@12 (37.3). In practical applications, the choice of r should be based on the desired balance between quality and robustness.

4.2 Human evaluation

While LLM-based evaluators serve as good proxies for human judgments in several cases (Zheng et al., 2023), their results should be interpreted with cau-

Metric →	Soft Win Rate					
Method →	KGW	Unigram	EXP	EXP-Edit	Blackbox	POSTMARK@12
LLAMA-3-8B	37	17	23	49	45	74
LLAMA-3-8B-INST	52	52	59	57	55	68
MISTRAL-7B-INST	57	54	49	54	46	64
GPT-4	-	-	-	-	53	64

Table 6: Soft win rates computed based on the pairwise comparison evaluation with GPT-4-TURBO as the judge, measured over 100 pairs of unwatermarked and watermarked OpenGen outputs from various LLMs (first column). POSTMARK@12 outperforms all baselines.

Configuration	Soft Win Rate	TPR After Para.
POSTMARK@6	84	20.8
POSTMARK@8	79	28.2
POSTMARK@12	64	59.4
POSTMARK@15	67	61.9
POSTMARK@20	62	82.8
POSTMARK@30	55	98

Table 7: Quality-robustness trade-off. All soft win rates are averaged over 100 pairs of unwatermarked and watermarked texts judged by GPT-4-TURBO. All paraphrased TPR numbers at 1% FPR are computed over on 500 OpenGen instances.

tion, as they can be biased to certain aspects of the text such as length (Wang et al., 2023) or overlap between the generator and the judge model (Panickssery et al., 2024). Thus, we hire two annotators fluent in English and conduct two human annotation studies detailed below, addressing Q3 and Q4. More details on annotator qualifications, payment, and each annotation setup can be found in §L.

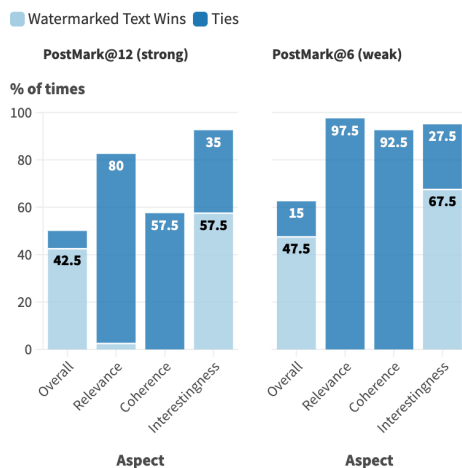


Figure 2: Pairwise preference human evaluation results on POSTMARK@12 and POSTMARK@6. For both configurations, the watermarked text is at least as good as its unwatermarked counterpart the majority of the time in all aspects.

> **Q3: POSTMARK watermarked texts are at least as good as their unwatermarked counterparts the majority of the time.** We first evaluate

the impact of POSTMARK on quality through a pairwise comparison task, similar to the setup in Section 4.1. Each annotator reads 20 OpenGen prefixes and the corresponding pairs of anonymized watermarked and unwatermarked responses generated by GPT-4. We then ask them to indicate their preferred response *overall*, as well as their preferences in terms of *relevance*, *coherence*, and *interestingness*, allowing for ties. Results in Figure 2 indicate that for POSTMARK@12 and POSTMARK@6, watermarked responses are at least as good as their unwatermarked counterparts the majority of the time (i.e., total percentage of wins and ties $\geq 50\%$). As expected, reducing the insertion rate to 6% improves quality, especially in the *coherence* aspect.¹¹ To put things in perspective, a previous human evaluation study by Kirchenbauer et al. (2024) found that annotators preferred KGW-watermarked text over unwatermarked text only 38.4% of the time.

> **Q4: Annotators struggle to identify the words inserted by POSTMARK.** A primary concern with POSTMARK is whether the words inserted into the watermarked text will be conspicuous enough for humans to identify, making it easy for attackers to remove them. To measure this, we create an anonymized mixture of 20 unwatermarked¹² and 20 watermarked responses generated for 20 prefixes in OpenGen with GPT-4 as the base LLM.¹³ We then ask annotators to highlight out-of-place words that they think might have been inserted post-hoc after the initial generation. Overall, annotators achieve a **F1** of merely **0.06** (0.46 precision, 0.03 recall). On average, they highlight 2.2 words in each unwatermarked response, and 3.45 words in each watermarked response. Thus, even when annotators are aware of the insertion of words, they cannot pinpoint the specific words.

5 Related work

Early research on watermarking: Our work is relevant to early work on watermarking text documents, either using the text document image (Brasnil et al., 1995; Low et al., 1998), syntactic transformations (Atallah et al., 2001; Meral et al., 2009),

¹¹While soft win rates computed from human annotations are much lower than those from GPT-4-TURBO’s judgments, both judges agree that a smaller r improves quality.

¹²We include unwatermarked responses in this evaluation as a baseline. For fairness, we regenerated unwatermarked texts to roughly match the length of the watermarked texts.

¹³These 20 prefixes are different from the ones they see in the pairwise comparison evaluation.

or semantic changes (Atallah et al., 2003; Topkara et al., 2006). Later work also explores watermarking machine-generated text (Venugopal et al., 2011).

Watermarking LLM-generated text: Recent research has primarily focused on watermarking LLM-generated outputs. Most existing approaches operate in the *whitebox* setting, assuming access to model logits and the ability to modify the decoding process (Fang et al., 2017; Kaptchuk et al., 2021; Aaronson and Kirchner, 2022; Kirchenbauer et al., 2023; Zhao et al., 2023; Liu et al., 2024a,b) or inject detectable signals without altering the original token distribution (Christ et al., 2023; Kuditipudi et al., 2024). Alternatively, Hou et al. (2023, 2024) watermark at the sentence level via rejection sampling. Prior *blackbox* methods access only model outputs (like POSTMARK), but rely on simple lexical substitution (Abdelnabi and Fritz, 2021; Qiang et al., 2023; Yang et al., 2023; Munyer et al., 2024).

Evading watermark detection: Our work also relates to prior work on text editing attacks designed to evade watermark detection. He et al. (2024) propose a cross-lingual attack, while Kirchenbauer et al. (2024) studies a copy-paste attack that embeds watermarked text into a larger human-written document. Krishna et al. (2023) train a controllable paraphraser that allows for control over lexical and syntactic diversity. Sadasivan et al. (2024) design a recursive paraphrasing attack that repeatedly rewrites watermarked text. Similar to our work, several studies directly prompt an instruction-following LLM to paraphrase text (Zhao et al., 2023; Hou et al., 2023, 2024; Liu et al., 2024b; Kirchenbauer et al., 2024).

Quality-robustness trade-off: Relevant to our discussion in Section 4, several recent papers highlight the impact of watermarking on quality. In line with our conclusions, Singh and Zou (2023) and Molenda et al. (2024) both find that less robust watermarks tend to have less negative impact on text quality.

6 Conclusion

We propose POSTMARK, a novel watermarking approach that only requires access to the underlying model’s outputs, making it applicable by third-party entities to outputs from API providers. Through extensive experiments across eight baseline algorithms, five base LLMs, and three datasets,

we show that POSTMARK is more robust to paraphrasing attacks than existing methods. We conduct a human evaluation to show that words inserted by POSTMARK are not easily identifiable by humans. We further run comprehensive quality evaluations covering *coherence*, *relevance*, *interestingness*, and *factuality*, and find that POSTMARK preserves text quality relatively well. Future work could look into further optimizing each of the three modules in POSTMARK, evaluating POSTMARK on attacks other than paraphrasing, or making logit-based methods less entropy-dependent.

Limitations

In this section, we address the primary limitations of our work.

Other attacks: Our work focuses on evaluating robustness of various watermarking methods against paraphrasing attacks. However, there are many other interesting and practical attacks that we do not consider, such as the copy-paste attack and the recursive paraphrasing attack discussed in Section 5. We anticipate that POSTMARK will be less effective when the watermarked text is embedded in a larger human-written document or when it undergoes repeated paraphrasing, similar to other watermarking methods. We leave the exploration of these other types of attacks to future work.

Runtime and API costs: The POSTMARK implementation used in all our main experiments relies on closed-source models from OpenAI (TEXT-EMBEDDING-3-LARGE and GPT-4O). As a result, the runtime and costs of running POSTMARK are heavily dependent on the API provider. Our cost estimate in §G suggests that watermarking 100 tokens with the default POSTMARK@12 configuration costs around \$1.2 USD. However, the framework is highly flexible in terms of module selection. In fact, as demonstrated in Section 3.2, an open-source implementation can perform nearly as well as the closed-source version. We leave the optimization of open-source implementations of POSTMARK to future work.

Ethical considerations

Our human study was determined exempt by IRB review. All annotators have consented to the release of their annotations, and we ensured they were fairly compensated for their valuable contributions. Scientific artifacts are implemented for

their intended usage. The risks associated with our framework are no greater than those already present in the large language models it utilizes (Weidinger et al., 2021).

Acknowledgments

We extend special gratitude to the Upwork annotators for their hard work. This project was partially supported by awards IIS-2202506 and IIS-2312949 from the National Science Foundation (NSF).

References

- Scott Aaronson and Hendrik Kirchner. 2022. [Watermarking gpt outputs](#).
- Sahar Abdelnabi and Mario Fritz. 2021. [Adversarial watermarking transformer: Towards tracing text provenance with data hiding](#). *Preprint*, arXiv:2009.03015.
- AI@Meta. 2024. [Llama 3 model card](#).
- Mikhail J. Atallah, Victor Raskin, Michael Crogan, Christian Hempelmann, Florian Kerschbaum, Dina Mohamed, and Sanket Naik. 2001. Natural language watermarking: Design, analysis, and a proof-of-concept implementation. In *Information Hiding*, pages 185–200, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Mikhail J. Atallah, Victor Raskin, Christian F. Hempelmann, Mercan Karahan, Radu Sion, Umut Topkara, and Katrina E. Triezenberg. 2003. Natural language watermarking and tamperproofing. In *Information Hiding*, pages 196–212, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *Preprint*, arXiv:2204.05862.
- J.T. Brassil, S. Low, N.F. Maxemchuk, and L. O’Gorman. 1995. [Electronic marking and identification techniques to discourage document copying](#). *IEEE Journal on Selected Areas in Communications*, 13(8):1495–1504.
- Miranda Christ, Sam Gunn, and Or Zamir. 2023. [Undetectable watermarks for language models](#). *Preprint*, arXiv:2306.09194.
- Together Computer. 2023. [Redpajama: an open dataset for training large language models](#).
- Tina Fang, Martin Jaggi, and Katerina Argyraki. 2017. [Generating steganographic text with LSTMs](#). In *Proceedings of ACL 2017, Student Research Workshop*, pages 100–106, Vancouver, Canada. Association for Computational Linguistics.
- Zhiwei He, Binglin Zhou, Hongkun Hao, Aiwei Liu, Xing Wang, Zhaopeng Tu, Zhuosheng Zhang, and Rui Wang. 2024. [Can watermarks survive translation? on the cross-lingual consistency of text watermark for large language models](#). *Preprint*, arXiv:2402.14007.
- Abe Bohan* Hou, Jingyu* Zhang, Tianxing* He, Yung-Sung Chuang, Hongwei Wang, Lingfeng Shen, Benjamin Van Durme, Daniel Khoshabi, and Yulia Tsvetkov. 2023. [Semstamp: A semantic watermark with paraphrastic robustness for text generation](#). In *Annual Conference of the North American Chapter of the Association for Computational Linguistics*.
- Abe Bohan Hou, Jingyu Zhang, Yichen Wang, Daniel Khoshabi, and Tianxing He. 2024. [k-semstamp: A clustering-based semantic watermark for detection of machine-generated text](#). *Preprint*, arXiv:2402.11399.
- Zhengmian Hu, Lichang Chen, Xidong Wu, Yihan Wu, Hongyang Zhang, and Heng Huang. 2024. [Unbiased watermark for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L elio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth ee Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Gabriel Kaptchuk, Tushar M. Jois, Matthew Green, and Aviel D. Rubin. 2021. [Meteor: Cryptographically secure steganography for realistic distributions](#). In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, CCS ’21*, page 1529–1548, New York, NY, USA. Association for Computing Machinery.
- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. 2023. [A watermark for large language models](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 17061–17084. PMLR.
- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Manli Shu, Khalid Saifullah, Kezhi Kong, Kasun Fernando, Aniruddha Saha, Micah Goldblum, and Tom Goldstein. 2024. [On the reliability of watermarks for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Kalpesh Krishna, Yixiao Song, Marzena Karpinska, John Wieting, and Mohit Iyyer. 2023. [Paraphrasing](#)

- evades detectors of AI-generated text, but retrieval is an effective defense. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Rohith Kuditipudi, John Thickstun, Tatsunori Hashimoto, and Percy Liang. 2024. [Robust distortion-free watermarks for language models](#). *Preprint*, arXiv:2307.15593.
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hannaneh Hajishirzi. 2024. [Rewardbench: Evaluating reward models for language modeling](#). *Preprint*, arXiv:2403.13787.
- Aiwei Liu, Leyi Pan, Xuming Hu, Shuang Li, Lijie Wen, Irwin King, and Philip S. Yu. 2024a. [An unforgeable publicly verifiable watermark for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Aiwei Liu, Leyi Pan, Xuming Hu, Shiao Meng, and Lijie Wen. 2024b. [A semantic invariant robust watermark for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- S.H. Low, N.F. Maxemchuk, and A.M. Lapone. 1998. [Document identification for copyright protection using centroid detection](#). *IEEE Transactions on Communications*, 46(3):372–383.
- Hasan Mesut Meral, Bülent Sankur, A. Sumru Özsoy, Tunga Güngör, and Emre Sevinç. 2009. [Natural language watermarking via morphosyntactic alterations](#). *Computer Speech and Language*, 23(1):107–125.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *International Conference on Learning Representations*.
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FActScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D. Manning, and Chelsea Finn. 2023. [Detectgpt: Zero-shot machine-generated text detection using probability curvature](#). *Preprint*, arXiv:2301.11305.
- Piotr Molenda, Adian Liusie, and Mark J. F. Gales. 2024. [Waterjudge: Quality-detection trade-off when watermarking large language models](#). *Preprint*, arXiv:2403.19548.
- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2023. [MTEB: Massive text embedding benchmark](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2014–2037, Dubrovnik, Croatia. Association for Computational Linguistics.
- Travis Munyer, Abdullah Tanvir, Arjon Das, and Xin Zhong. 2024. [Deeptextmark: A deep learning-driven text watermarking approach for identifying large language model generated text](#). *Preprint*, arXiv:2305.05773.
- Roberto Navigli, Simone Conia, and Björn Ross. 2023. [Biases in large language models: Origins, inventory, and discussion](#). *J. Data and Information Quality*, 15(2).
- Zach Nussbaum, John X. Morris, Brandon Duderstadt, and Andriy Mulyar. 2024. [Nomic embed: Training a reproducible long context text embedder](#). *Preprint*, arXiv:2402.01613.
- OpenAI. [Model release blog: GPT-4o](#). Technical report, OpenAI.
- OpenAI. 2024a. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- OpenAI. 2024b. [New embedding models and api updates](#).
- Leyi Pan, Aiwei Liu, Zhiwei He, Zitian Gao, Xuandong Zhao, Yijian Lu, Binglin Zhou, Shuliang Liu, Xuming Hu, Lijie Wen, and Irwin King. 2024. [Mark-llm: An open-source toolkit for llm watermarking](#). *Preprint*, arXiv:2405.10051.
- Arjun Panickssery, Samuel R. Bowman, and Shi Feng. 2024. [Llm evaluators recognize and favor their own generations](#). *Preprint*, arXiv:2404.13076.
- Jipeng Qiang, Shiyu Zhu, Yun Li, Yi Zhu, Yunhao Yuan, and Xindong Wu. 2023. [Natural language watermarking via paraphraser-based lexical substitution](#). *Artif. Intell.*, 317(C).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1).
- Vinu Sankar Sadasivan, Aounon Kumar, Sriram Balasubramanian, Wenxiao Wang, and Soheil Feizi. 2024. [Can ai-generated text be reliably detected?](#) *Preprint*, arXiv:2303.11156.
- Ilya Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson. 2023. [The curse of recursion: Training on generated data makes models forget](#). *Preprint*, arXiv:2305.17493.
- Karanpartap Singh and James Zou. 2023. [New evaluation metrics capture quality degradation due to llm watermarking](#). *Preprint*, arXiv:2312.02382.
- Edward Tian. 2023. [Gptzero: An ai text detector](#).

- Umüt Topkara, Mercan Topkara, and Mikhail J. Atallah. 2006. [The hiding virtues of ambiguity: quantifiably resilient watermarking of natural language text through synonym substitutions](#). In *Proceedings of the 8th Workshop on Multimedia and Security*, page 164–174, New York, NY, USA. Association for Computing Machinery.
- Ashish Venugopal, Jakob Uszkoreit, David Talbot, Franz Och, and Juri Ganitkevitch. 2011. [Watermarking the outputs of structured prediction with an application in statistical machine translation](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1363–1372, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023. [Large language models are not fair evaluators](#). *Preprint*, arXiv:2305.17926.
- Yequan Wang, Jiawen Deng, Aixin Sun, and Xuying Meng. 2022. [Perplexity from plm is unreliable for evaluating text quality](#).
- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, Zac Kenton, Sasha Brown, Will Hawkins, Tom Stepleton, Courtney Biles, Abeba Birhane, Julia Haas, Laura Rimell, Lisa Anne Hendricks, William Isaac, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2021. [Ethical and social risks of harm from language models](#). *Preprint*, arXiv:2112.04359.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2015. [From paraphrase database to compositional paraphrase model and back](#). *Transactions of the Association for Computational Linguistics*, 3:345–358.
- Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. 2024. [Hallucination is inevitable: An innate limitation of large language models](#). *Preprint*, arXiv:2401.11817.
- Xi Yang, Kejiang Chen, Weiming Zhang, Chang Liu, Yuang Qi, Jie Zhang, Han Fang, and Nenghai Yu. 2023. [Watermarking text generated by black-box language models](#). *Preprint*, arXiv:2305.08883.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. [Opt: Open pre-trained transformer language models](#). *Preprint*, arXiv:2205.01068.
- Xuandong Zhao, Prabhanjan Ananth, Lei Li, and Yu-Xiang Wang. 2023. [Provable robust watermarking for ai-generated text](#). *Preprint*, arXiv:2306.17439.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhaghao Wu, Yonghao Zhuang, Zi Lin,

A More details on the vocabulary \mathbb{V} of the SECTABLE

In this section, we provide more details on the creation of SECTABLE, and address how often a word in the SECTABLE can be selected as a watermark word.

Filtering the SECTABLE vocabulary \mathbb{V} : Specifically, we restrict \mathbb{V} to only include lowercase nouns, verbs, adjectives, and adverbs that occur at least 1,000 times in the WikiText-103 training split. This results in a final vocabulary of 3,266 words.

Frequency of words chosen as watermark words: In Figure 3, we plot the frequency distribution of all watermark words obtained for 500 OpenGen outputs (generated with GPT-4 as the base LLM). We find that the majority of the words are only selected as watermark words for less than 5% of all outputs, while two major hub words are selected in more than 20% of the outputs. Overall, the hubness problem is not too severe, but it could be mitigated by a more careful selection of the embeddings used in the SECTABLE.

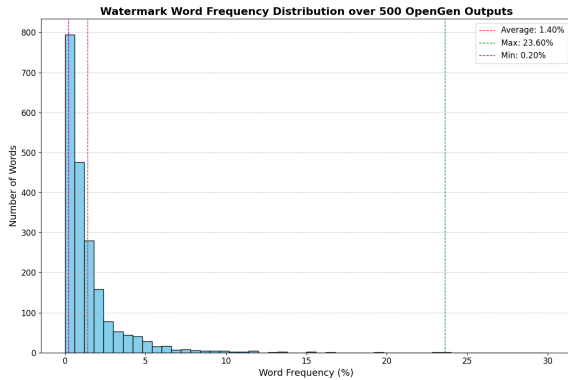


Figure 3: Watermark word frequency distribution over 500 OpenGen outputs. The majority of the words are chosen as watermark words less than 5% of the time. There are only two major hub words that are selected more than 20% of the time.

B Prompt for the INSERTER

Given below are a piece of text and a word list. Rewrite the text to incorporate all words from the provided word list. The rewritten text must be coherent and factual. Distribute the words from the list evenly throughout the text, rather than clustering them in a single

	SIM(positive)	SIM(negative)
PARAGRAM	64.8	2.4
GLOVE	60.7	16.4
NOMIC-EMBED	59.9	33.2
TEXT-EMBEDDING-3-LARGE	64.2	29.8

Table 8: Cosine similarity between embeddings of positive pairs (word + its synonym) and between negative pairs (word + irrelevant word) computed with different embedding models, averaged over 174 tuples of (word, synonym, irrelevant word).

section. When rewriting the text, try your best to minimize text length increase. Only return the rewritten text in your response, do not say anything else.

Text: {}

Word list: {}

Rewritten text:

C More details on cosine similarity word matching during detection

We use the PARAGRAM word embedding model developed by Wieting et al. (2015) to perform cosine similarity word matching during detection. We find this model to be superior at distinguishing semantically related words from irrelevant words, see details in Table 8.

D More details on baselines

In this section, we provide more details on how we run our baselines.

D.1 Expanded descriptions of baselines

(1) **KGW** (Kirchenbauer et al., 2023): Partitions the vocabulary into “green” and “red” lists based on the previous token, then boosts the probability of green tokens during generation. Detection is done by comparing the number of green tokens present to the expected count under the null hypothesis of no watermarking. (2) **Unigram** (Zhao et al., 2023): A variant of KGW that uses a fixed green-red partition for all generations instead of re-partitioning the vocabulary at each token, making it more robust to editing attacks. (3) **EXP** (Aaronson and Kirchner, 2022): Uses exponential sampling to embed a watermark by biasing token selection with a pseudo-random sequence during text generation. Detection measures the correlation between the generated text and the sequence to identify

the watermark. (4) **EXP-Edit** (Kuditipudi et al., 2024): A variant of the EXP watermark that incorporates edit distance to measure the correlation. (5) **SemStamp** (Hou et al., 2023): A sentence-level algorithm that partitions the semantic space using locality-sensitive hashing with arbitrary hyperplanes, assigning binary signatures to regions and accepting sentences that fall within “valid” regions, which enhances robustness against paraphrase attacks. (6) **k-SemStamp** (Hou et al., 2024): Improves upon SemStamp by using k-means clustering to partition the semantic space. (7) **SIR** (Liu et al., 2024b): Generates watermark logits from the semantic embeddings of preceding tokens using an embedding language model and a trained watermark model. These logits are added to the language model’s logits. Detection works by averaging these watermark logits for each token and identifying a watermark if the average is significantly greater than zero. (8) **Blackbox** (Yang et al., 2023): While all other baseline methods require access to model logits, this method focuses on the blackbox setting where only the model output is observable, similar to our assumption. It encodes words as binary bits, replaces bit-0 words with synonyms representing bit-1, and detects watermarks through a statistical test identifying the altered distribution of binary bits.

D.2 Hyperparameters for baselines

All baselines are run with nucleus sampling with $p = 0.9$ unless otherwise specified.

KGW: We run KGW in the LeftHash configuration with $\gamma = 0.5$ and $\delta = 4.0$, using the original authors’ implementation. These hyper-parameters control the size of the green token list and the strength of the watermark, respectively. While δ is typically set to 2.0 in prior literature, we chose $\delta = 4.0$ based on findings by Kirchenbauer et al. (2024). They found that $\delta = 4.0$ made the watermark more robust to paraphrasing attacks in their experiments with Vicuna, a supervised instruction-finetuned model. Given that our experiments also focus on lower-entropy models aligned through RLHF or instruction tuning, we adopt the same value for δ .

Unigram: To align with the setup of KGW, we set $\gamma = 0.5$ and $\delta = 4.0$ for Unigram as well. While the authors open-source their code, we ran into unexpected performance issues, where Unigram could not even achieve a TPR at 1%

FPR higher than 70% even before any attacks on OpenGen with LLAMA-3-8B as the base model. Thus, we switched to the implementation in MarkLLM (Pan et al., 2024), an open-source watermarking toolkit. With this implementation, Unigram’s TPR before attacks became close to 100% and the TPR after attacks stayed above 90%, in line with results reported in the Unigram paper (Zhao et al., 2023).

EXP: We run EXP with prefix length set to 1 using the MarkLLM implementation.

EXP-Edit: Using the authors’ implementation, we run EXP-Edit with $\gamma = 0.5$, watermark key length = 256, block size = sequence length = 300, and number of resamples = 100. This method is run with multinomial sampling (the default setting in the authors’ code), because we find that adding a nucleus sampling logits wrapper on top significantly hurts its performance. For LLAMA-3-8B-INST and MISTRAL-7B-INST, we find that this method cannot reach a TPR at 1% FPR above 70% even before attacks. We tried several values for γ , the hyperparameter that controls the statistical power of the watermark, but it did not improve the results. Increasing the number of resamples to 500 also had little effect.

Blackbox: We run Blackbox with $\tau = 0.8$ and $\lambda = 0.83$ using fast detection with the authors’ implementation. Empirically, we find that fast detection offers a significant speed advantage with negligible impact on performance when compared to precise detection. On 200 OpenGen outputs with GPT-4 as the base LLM, using precise detection yields TPR of 100 before paraphrasing and 3.5 after paraphrasing, whereas fast detection yields 99 and 0.5.

E More details on base models

In this section, we provide more details on how we run the base generator models.

Model checkpoints: We detail the checkpoint we use for each base model in Table 9.

Generation length: For all aligned models (LLAMA-3-8B-INST, MISTRAL-7B-INST, and GPT-4), we generate free-form text until the model outputs an EOS (end-of-sequence) token to simulate the downstream setting. For LLAMA-3-8B, we set the maximum token limit to 300, as generating

Model	Checkpoint
LLAMA-3-8B	Meta-Llama-3-8B
LLAMA-3-8B-INST	Meta-Llama-3-8B-Instruct
MISTRAL-7B-INST	Mistral-7B-Instruct-v0.2
GPT-4	OpenAI API (gpt-4-0613)

Table 9: Base model checkpoints.

freely until reaching EOS often leads to meaningless repetitions, sometimes even exceeding 8,000 tokens. We do not run OPT-1.3B ourselves.

F Paraphrasing attack setup

In this section, we provide more details on the paraphrasing attack we use for all experiments.

Prompt for sentence-level paraphrasing: We build on the prompt used by Hou et al. (2023, 2024) and include more clarification on what to return:

Given some previous context and a sentence following that context, paraphrase the current sentence. Only return the paraphrased sentence in your response.

Previous context: {}

Current sentence to paraphrase: {}

Your paraphrase of the current sentence:

Why sentence-level paraphrasing? We choose a sentence-level paraphrasing setup for two reasons. First, Hou et al. (2023, 2024) use a sentence-level paraphrasing setup to evaluate the robustness of their method. Since we are unable to run their method directly, adopting the same paraphrasing setup allows for a fair comparison with their results. Second, as observed by Kirchenbauer et al. (2024), naively prompting GPT-3.5-TURBO to rewrite the entire input text often results in significant loss of important content. While the authors developed a sophisticated prompt to mitigate this issue, we empirically find that paraphrasing at a sentence level achieves a similar effect.

G POSTMARK runtime and API cost estimates

Runtime: We compare the runtime of several POSTMARK configurations with other baselines in Table 10. Recall that in our experiments, we find insertion success rate to be higher if we divide the watermark word list into sublists of 10 words, then ask the INSERTER to insert one sublist at a time. This iterative insertion process can have some

Method	Avg Time / Output
POSTMARK@6	29.2
POSTMARK@12	36.6
POSTMARK@12 (no iter.)	25.3
KGW	17.5
Unigram	18.5
EXP	18.4
EXP-Edit	17.3
Blackbox	21.6

Table 10: Average time (in seconds) it takes to generate one watermarked instance with LLAMA-3-8B-INST as the base LLM. Runtime is averaged over 10 outputs, with an average token count of 280. For POSTMARK and Blackbox, the runtime includes the time it takes for LLAMA-3-8B-INST to generate the initial unwatermarked output. POSTMARK@12 (no iter.) refers to the setup where instead of breaking up the watermark word list into sublists and iteratively asking the INSERTER to insert one sublist at a time, we directly ask the INSERTER to insert all words in the list.

negative impact on runtime, but it may become unnecessary in the future when the INSERTER has better instruction-following capabilities.

API costs: Under the default POSTMARK@12 configuration with GPT-4O as the INSERTER and TEXT-EMBEDDING-3-LARGE as the EMBEDDER watermarking 500 outputs with around 300 tokens costs around \$18.5 USD, which means that watermarking 100 tokens costs about \$1.2 on average.

H POSTMARK length comparison

We present a comparison between output length (before and after watermarking) for various watermarking methods in Table 11.

I Unigram repetitions

We present several examples of Unigram’s repetitive watermarked outputs in Table 12, generated with LLAMA-3-8B as the base LLM.

J EXP repetitions

We present several examples of EXP’s repetitive watermarked outputs in Table 13, generated with OPT-1.3B as the base LLM.

K Prompt for the LLM-based pairwise evaluation setup

Please act as an impartial judge and evaluate the quality of the text completions provided by two large language models to the prefix displayed

Metric → Methods →	Number of Tokens (Before / After Watermarking)					
	KGW	Unigram	EXP	EXP-Edit	Blackbox	POSTMARK@12
LLAMA-3-8B	239.6 / 226.6	237.6 / 250.7	232.5 / 269.8	213 / 225.7	239.6 / 244.8	239.6 / 381.2
LLAMA-3-8B-INST	251.2 / 234.6	259.5 / 261.6	259 / 282.6	251.3 / 255	251.2 / 256.4	251.2 / 431
MISTRAL-7B-INST	315.3 / 588.2	318 / 321	317.4 / 247.8	248.7 / 249.5	315.3 / 320.6	315.3 / 552.2
GPT-4	-	-	-	-	301.2 / 305.7	301.2 / 507.1

Table 11: Length comparison between different watermarking methods before and after watermarking, averaged over 500 OpenGen outputs.

below. Assess each response according to the criteria outlined. After scoring each criterion, provide a summary of your evaluation for each response, including examples that influenced your scoring. Additionally, ensure that the order in which the responses are presented does not affect your decision. Do not allow the length of the responses to influence your evaluation. Be as objective as possible.

Criteria:

1. Relevance to the prefix
2. Coherence
3. Interestingness

Start with a brief statement about which response you think is better overall. Then, for each criterion, state which response is better, or if there is a tie, followed by a concise justification for that judgment. At the very end of your response, declare your verdict by choosing one of the choices below, strictly following the given format: "[[A]]" if assistant A is better overall, "[[B]]" if assistant B is better overall, or "[[C]]" for a tie.

[Prefix]

{}

[Response A]

{}

[Response B]

{}

L Human evaluation setup and costs

Hiring annotators: We hire two annotators from [Upwork](#). Both annotators are fluent in English,

have 100% job success rates, and have demonstrated exceptional professionalism in their communications with us.

Pairwise evaluation: The interface we use for this task, built with [Label Studio](#), is shown in [Figure 4](#). For this task, we pay each annotator \$2 USD per pair, and they spend around 5-10 minutes per pair.

Identifying watermark words: The interface we use for this task is shown in [Figure 5](#). For this task, we pay each annotator \$1.5 USD per output, and they spend around 3-5 minutes on each output.

Prefix	Without Watermarking	After Unigram Watermarking
<p>Unlike mountains of similar altitude elsewhere, Elbert lacks both a permanent snowpack and a prominent north-facing cirque, which can be attributed to its position among other mountains of similar height, causing it to receive relatively small quantities of precipitation. Mount Elbert was named by miners in honor of Samuel Hitt Elbert, the governor of the then-Territory of Colorado, because he brokered a treaty in September 1873 with the Ute tribe that opened up more than of reservation land to mining and railroad activity.</p>	<p>Mount Elbert Mountain, United States of America: Elevation: 4401 meters/ 14434 feet: USGS Map Name: Mountain Elbert: Latitude: 39° 07' 17" N Longitude: 106° 26' 08" W Latitude/Longitude (WGS84) 39° 7' 17" N 106° 26' 8" W (39.1213939, -106.4355046) Elevation: 4400+ft. Name Coordinates: Elevation Elbert, Mount: 39° 07' 17" N 106° 26' 08" W: 4,401 m (14,440 ft) Location: Colorado, United States: Parent peak: Uncompahgre Peak: Elevation gain: 4,419 ft (1,348 m) Coordinates: 39° 07' 17" N, 106° 26' 08" W; OS grid reference: Coordinates: Denomination: From the list below, please choose the region you would like to visit Elbert, Mount, and we'll show you the times of departure and scheduled arrival, in addition to realistic information about guarantees, time change schedules, and delays for this train. Mount Elbert, the highest mountain in Colorado, stands at 4,401m.</p>	<p>Mount of Snow – Mt. Mt. Canyon Lake High School alumnus Steve Phelps of Leadville competes annually on Mount. The summit of Mt. Because of its high altitude of approximately Mt. Mt. Hiking Mt. Mounts of Snow – Mount of Snow is peak of Mount. Mt. Mount of Snow – Mount of Snow is peak of Mount. Holly Berry For. Mount of Snow – Mount of Snow is peak of Mount. Mounts of Snow – Mount of Snow is peak of Mount. Mounts of Snow – Mount of Snow is peak of Mount. Mounts of Snow – Mount of Snow is peak of Mount. Mounts of Snow – Mount of Snow is peak of Mount. Mounts of Snow – Mount of Snow is peak of Mount. Mount of Snow – Mount of Snow is peak of Mount. Mount of Snow – Mount of Snow is peak of Mount. Mount of Snow – Mount of Snow is peak of Mount. Mount of Snow – Mount of Snow is peak of Mount. Mount of Snow – Mount of Snow is peak of Mount. Mount of Snow – Mount of Snow is peak of Mount. Mt. Mt. Mount of Snow – Mount of Snow is peak of Mount. Holly Berry For. Mt. Mount of Snow – Mount of Snow is peak of Mount. Mount of Snow – Mount of Snow is peak of Mount. Mount of Snow – Mount of Snow is peak of Mount. Mt. Mount of Snow – Mount of Snow is peak of Mount.</p>
<p>We are not! Like men in the story of the Good Samaritan, they pass by on the other side...</p>	<p>they avoid the God-driven endeavors of the church. Many modernized men avoid the stories of the women in the Bible and opt to ignore Jesus' directive to be "transformed into his image," Colossians 3:10. (Not the way our men view the world through "TrophyWives Corner"!). But, this is certainly not all. Now, as of today, we will work with a modernized spin on God's word to show you how Jesus, 2000 years ago, managed to "reformat" most men through his efforts in church. But before we start, I must thank two lovely women who were in my life in the 1970's...women who unknowingly were reformatting the world of a confused young boy like me. I really appreciate you both, no matter how confusing you made the world seem to me then. And by the way, you will meet a couple of young lads too. I'm sure your daughters, granddaughters and nieces will relate to this. This is for all the little girls out there. Thank you guys, and keep on telling us men, just like Jesus did, what we should be like.</p>	<p>of Route 119. But that leaves you to cross Route 119 without adult supervision. Presuming you've had adult supervision all along. Which is probably true if you haven't had all of that adult supervision that ought to have been all along. The point is that you're here to cross Route 119 without adult supervision. Like the Good Fisherman who nets all of us into Heaven. The only catch is that you probably won't be able to cross without being hurt. The flip side is that you can probably cross without hurting anyone else. If you give to God without doubting that you'll get what you give to him then you'll be able to cross Route 119 without adult supervision. Giving to God is almost like going to Heaven because what we give to him is what we'll get back. If you give him bread then you'll get bread. If you give him food then you'll get food. If you give him money then you'll get money. If you give him parents then you'll get parents. If you give him teachers then you'll get teachers. If you give him insurance then you'll get insurance. If you give him Good Parents then you'll get Good Parents. If you give him Good Men then you'll get Good Men. If you give him Good Fisherman then you'll get Good Fisherman. If you give him Good Fish then you'll get Good Fish. If you give him Good Charismata then you'll get Good Charismata.</p>

Table 12: Example repetitive outputs by Unigram with LLAMA-3-8B as the base LLM.

Prefix	Without Watermarking	After Unigram Watermarking
Cluster comprises IBM's Opteron-based eServer 325 server and systems management software and storage devices that can run Linux and Windows operating systems.	It also supports Oracle, MS-Access, and B2Vista. The server runs clustered versions of Unix. It supports Virtual SAN. It has one onboard hard drive and two externally attached hard drives. The server has built-in Web browser and FTP server capabilities. It comes with a Linux-based telnet interface. The user interface is based on the Konsole language. The software supports file and directory browsing. It can be configured to run under different operating system versions. It supports file transfers. It supports file copying and editing. It supports backup and recovery. The software also supports network connectivity and file sharing. The user interface and documentation are in the Linux project's Make it Web (http://www.linux.org/makeitweb/). The eServer 325 supports hardware and software RAID. It supports mirrored drives and RAID 5. It supports RAID 3, HEAF, and RAID 5. It supports RAID 5's identical mirrored copies. Cluster servers can support multiple users.	This cluster is an addition to the eServer Performance System from IASC, which contains underlying servers running both Linux and Windows. This cluster is an addition to the eServer Performance System from IASC, which contains underlying servers running both Linux and Windows. This cluster is an addition to the eServer Performance System from IASC, which contains underlying servers running both Linux and Windows. This cluster is an addition to the eServer Performance System from IASC, which contains underlying servers running both Linux and Windows. This cluster is an addition to the eServer Performance System from IASC, which contains underlying servers running both Linux and Windows. This cluster is an addition to the eServer Performance System from IASC, which contains underlying servers running both Linux and Windows. This cluster is an addition to the eServer Performance System from IASC, which contains underlying servers running both Linux and Windows. This cluster is an addition to the eServer Performance System from IASC, which contains underlying servers running both Linux and Windows. This cluster is an addition to the eServer Performance System from IASC, which contains underlying servers running both Linux and Windows.
An eagerly awaited report from the World Health Organization (WHO) states that processed meats such as bacon and sausages cause cancer, and red meat likely does so too.	An eagerly awaited report from the World Health Organization (WHO) states that processed meats such as bacon and sausages cause cancer, and red meat likely does so too. The report is based on data from 44 countries. Two-thirds of the cancers occurred in women who ate red meat. Three-quarters of the cancers were in the male-only population. Four-fifths of the cancers were in the female-only population. In addition to bacon and sausages, high-fat meats such as cream and butter are also linked to cancer. "This is a message which has been circulating for a long time," said Dr. Margaret Hodge, head of the WHO's health department. She cited evidence that the link goes beyond processed meats. It is also linked to other risk factors, including poor nutrition, smoking, and obesity. She said people can make good choices when it comes to food. For example, they can limit consumption of red meat, but not processed meat. Processed meats are the most common type of meat in shops and markets around the world. The WHO's report lists the top 10 countries with the highest percentage of women who eat red meat.	This statement doesn't make mass damage to the country's meat supply quite possible, but the country's meat and agriculture industries have been decimated. This statement doesn't make mass damage to the country's meat supply quite possible, but the country's meat and agriculture industries have been decimated. This statement doesn't make mass damage to the country's meat supply quite possible, but the country's meat and agriculture industries have been decimated. This statement doesn't make mass damage to the country's meat supply quite possible, but the country's meat and agriculture industries have been decimated. This statement doesn't make mass damage to the country's meat supply quite possible, but the country's meat and agriculture industries have been decimated. This statement doesn't make mass damage to the country's meat supply quite possible, but the country's meat and agriculture industries have been decimated. This statement doesn't make mass damage to the country's meat supply quite possible, but the country's meat and agriculture industries have been decimated. This statement doesn't make mass damage to the country's meat supply quite possible, but the country's meat and agriculture industries have been decimated. This statement doesn't make mass damage to the country's meat supply quite possible, but the country's meat and agriculture industries have been decimated. This statement doesn't make mass damage to the country's meat supply quite possible, but the country's meat and agriculture industries have been decimated. This statement doesn't make mass damage to the country's meat supply quite possible, but the country's meat and agriculture industries have been decimated.

Table 13: Example repetitive outputs by EXP with OPT-1.3B as the base LLM.

Prefix

1 One tank tumbled down an embankment into the Tenaru River, drowning its crew. At 23:00 on 14 September, the remnants of the Kuma battalion conducted another attack on the same portion of the Marine lines, but were repulsed.

A

1 Despite the grim situation, the Kuma battalion demonstrated their tenacious resolve. Made up of the survivors of the earlier onslaught and bolstered by additional reinforcements, the battalion regrouped and launched a renewed assault as midnight cloaked the already intense battlefield. However, their initial attack under the cover of darkness proved futile once again against the iron-clad defense of the Marine lines.

2

3 Facing the grim odds, the commander of the Kuma battalion, Major General Masao, decided on a daring strategic maneuver. Instead of another focused assault on the same portion of the Marine lines, he ordered his troops to perform a flanking maneuver. The troops used the dense jungle as their cover, aiming to hit the Marine's defenses from the side and rear.

4

5 Regardless of clever strategy and audacity, the Kuma battalion was up against a formidable foe. The Marines, who had anticipated such a maneuver, were ready for them. As the Japanese troops broke from the growth, they were met with a hailstorm of bullets and artillery fire. Shouts and cries were drowned by the incessant bombardment. Within minutes, the jungle floor was carpeted with the fallen bodies of the would-be attackers from Kuma battalion.

6

7 By daybreak, it had become apparent that the Kuma battalion had been effectively wiped out. The few survivors were either captured or receded back into the jungle, leaving behind nothing but the haunting memories of a battlefield seeped in valor, and awash with chilling chaos. Their second attack on the Marine lines had ended in much the same way as the first - in dismal failure.

B

1 Despite the grim situation, the Kuma battalion demonstrated their tenacious resolve. Made up of the survivors of the earlier onslaught and bolstered by additional reinforcements, the battalion regrouped and launched a renewed assault as midnight cloaked the already intense battlefield. However, their initial attack under the cover of darkness proved futile once again against the iron-clad defense of the Marine regiment.

2

3 Facing the grim odds, the commander of the Kuma battalion, Major General Masao, decided on a daring strategic maneuver. Instead of another focused assault on the same portion of the Marine lines, he ordered his troops to perform a flanking maneuver. The troops used the dense jungle as their cover, aiming to hit the Marine's defenses from the side and rear.

4

5 Regardless of clever strategy and audacity, the Kuma battalion was up against a formidable foe. The Marines, who had anticipated such a maneuver, were ready for them. As the Japanese troops broke from the growth, they were met with a hailstorm of bullets and artillery fire. Shouts and cries were drowned by the incessant bombardment. Within minutes, the jungle floor was carpeted with the fallen bodies of the would-be attackers from Kuma battalion.

6

7 By daybreak, it had become apparent that the Kuma battalion had been effectively wiped out. The few survivors were either captured or receded back into the jungle, leaving behind nothing but the haunting memories of a battlefield seeped in valor, and awash with chilling chaos. Their second attack on the Marine lines had ended in much the same way as the first - in total dismal failure.

8

9 The magnitude of the battle was such that it drew the attention of multiple journalists, who documented the historical event. Major General Masao, an influential man, was noted for his protective stance towards his troops, despite the dismal outcome. The collaboration among the Marines proved to be a decisive factor, and the heroism displayed on both sides was undeniable. The battle remains a poignant chapter in military history, illustrating the relentless spirit and the harsh realities of war. The rise and fall of the Kuma battalion resemble the tragic cycles often seen in such conflicts, and the need to remake strategies in the face of overwhelming odds.

Overall, which text do you prefer as a continuation of the prefix?

A^[1] B^[2] Tie^[3]

Which text is more relevant to the prefix?

A^[4] B^[5] Tie^[6]

Which text is more coherent (i.e. has better logical flow)?

A^[7] B^[8] Tie^[9]

Which text is more interesting to read?

A^[10] B^[11] Tie^[12]

Figure 4: Human annotation interface for the pairwise comparison task.

weird 1

1 The RAAF's commitment in this operation was a testament to their adaptability and resilience in the face of a changing global conflict landscape. The missions marked the introduction of precise targeting and minimal collateral damage, signalling a significant shift in the realm of modern warfare, closely aligning with new military doctrine. This historical shift was of grand magnitude, reflecting the evolution from their predecessors' methods and the rise of new strategies.

2

3 The Squadron's deployment of the laser-guided bombs was seen as an evolution of combat strategy, replacing the traditional carpet-bombing approach with a more surgical and precise tactic. This not only reflected advances in technology but also mirrored the increasing global awareness of the need to limit civilian casualties in conflict situations, an aim that was deemed appropriate and necessary. The protective measures taken were a race against time to ensure lesser harm to civilians, a task that was handled with utmost care.

4

5 The implications of this change were profound. The more accurate approach meant that our forces were able to disrupt the enemy's infrastructure and resources without causing widespread devastation. It also meant that missions were safer for our pilots, as the precision bombs reduced the need for flying at dangerously low altitudes to ensure target accuracy, thus minimizing the risk of false targeting. The pilots who flew these missions were seen as heroes for their bravery and skill, having undergone rigorous training to master this new type of warfare.

6

7 The success of the 350 sorties flown by No. 75 Squadron during the invasion was a testament to the crews' training, skill, and tenacity. The sortie rate was relentless, but the Squadron maintained high operational integrity and effectiveness, managing to achieve their strategic objectives with minimum collateral damage, demonstrating their dominance in the air. Their protective actions ensured the safety of both military and civilian lives, a total success in their mission.

8

9 The deployment was not just about bomb runs. It also involved multiple roles such as territorial defence, aerial reconnaissance, and close air support, demonstrating the versatility and adaptability of both the RAAF and No. 75 Squadron. The success of these missions highlighted the evolved capabilities of the RAAF, and marked a new era in military operations, achieved through close collaboration and strategic planning. The attempt to integrate various roles was as harmonious as a bass line in a well-composed symphony, showcasing the grand scale of their achievements. Shortly after the operation, a survey was conducted to assess the effectiveness of the new strategies, and the results were overwhelmingly positive, despite some tasks being unsuccessfully completed. The regiment's efforts were akin to a well-fought tournament, where every move was calculated and executed with precision.

Figure 5: Human annotation interface for the watermark word identification task.