

# Is Multilingual LLM Watermarking Truly Multilingual? Scaling Robustness to 100+ Languages via Back-Translation

Anonymous Authors<sup>1</sup>

## Abstract

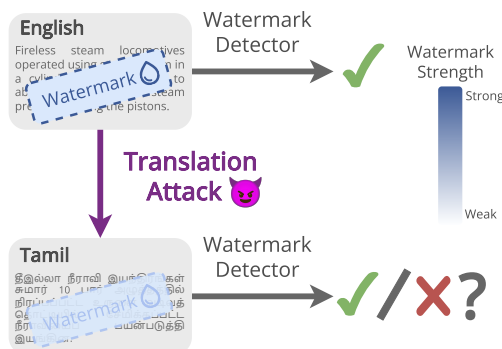
Multilingual watermarking aims to make large language model (LLM) outputs traceable across languages, yet current methods still fall short. Despite claims of cross-lingual robustness, they are evaluated only on high-resource languages. We show that existing multilingual watermarking methods are not truly multilingual: they fail to remain robust under translation attacks in medium- and low-resource languages. We trace this failure to semantic clustering, which fails when the tokenizer vocabulary contains too few full-word tokens for a given language. To address this, we introduce STEAM, a detection method that uses Bayesian optimisation to search among 133 candidate languages for the back-translation that best recovers the watermark strength. It is compatible with any watermarking method, robust across different tokenizers and languages, non-invasive, and easily extendable to new languages. With average gains of +0.23 AUC and +37%p TPR@1%, STEAM provides a scalable approach toward fairer watermarking across the diversity of languages.

## 1. Introduction

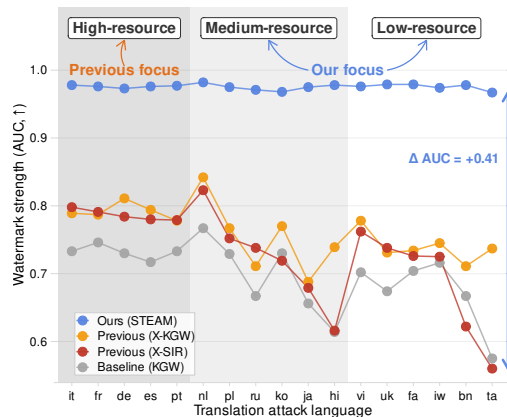
Recent advances in multilingual watermarking claim to make large language model (LLM) outputs traceable across languages. Yet existing methods have been evaluated only on a small set of high-resource languages, leaving open the question of whether these techniques truly generalise to the world’s linguistic diversity. In this work, we show that *current multilingual watermarking methods are not truly multilingual*. Their robustness weakens considerably for medium- and low-resource languages, revealing a major gap in current approaches to content provenance.

<sup>1</sup>Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.



(a) Translation attack against LLM watermarking.



(b) Watermarking robustness across languages.

Figure 1. (a) **Our goal** is to evaluate the robustness of LLM watermarks against translation attacks. (b) **Our analysis** reveals that existing multilingual watermarks fail to generalise across languages, while our approach (STEAM) performs consistently better across a wide range of languages overlooked by previous work.

The limited robustness of multilingual watermarking has broad consequences. Watermarking was designed to identify LLM-generated text and to reduce the spread of misinformation on social media and synthetic content on the web. An adversary can exploit *translation attacks*, in which a model generates text in one language and the content is translated into another, effectively scrubbing the watermark and reducing its strength (He et al., 2024; Al Ghanim et al., 2025; Han et al., 2025; Luo et al., 2025; Chen et al., 2025).

Figure 1a illustrates a translation attack. This threat is not theoretical: large-scale deployed systems such as Google’s SynthID (Dathathri et al., 2024), used in Gemini, Veo, Imagen, and others, lose detectability after translation (Han et al., 2025). This vulnerability could enable undetectable synthetic content to spread in hundreds of languages, particularly in communities where moderation tools are less effective.

We focus specifically on translation attacks because they are the natural threat in multilingual deployment, where LLM-generated content frequently spreads across language communities via translation. Other attacks, such as paraphrasing or sample-based mixing, operate within the same language and preserve the overall token distribution. Several existing methods already target this regime through semantic invariance (Zhao et al., 2023; Liu et al., 2024a; Ren et al., 2024). Translation attacks are different: they shift text to a different language, causing the token distribution to change substantially, which degrades the watermark signal in ways that semantic invariance cannot address.

*Semantic clustering* has been proposed as a multilingual extension of watermarking. It groups semantically equivalent tokens (for example, ‘house’, ‘maison’, ‘casa’) into clusters and treats all tokens in a cluster identically regarding the watermark key (for instance, all green or all red). While this approach performs adequately for high-resource languages, we observe that it performs poorly for many others. Tokenizers allocate tokens according to language frequency in their training data, meaning that only high-resource languages contain enough whole-word tokens to be properly represented in semantic clusters. For medium- and low-resource languages, most words are split into subword units not represented in any cluster, which substantially weakens the watermark. These findings suggest that semantic clustering cannot scale effectively beyond high-resource languages.

To address the limited robustness of semantic clustering, we introduce **STEAM**  $\hat{\mathcal{M}}$  (*Search-based Translation-Enhanced Approach for Multilingual watermarking*), a detection-time method that uses Bayesian optimisation to select the back-translation language that best recovers watermark strength. STEAM is non-invasive, model-agnostic, and compatible with any existing watermarking technique and tokenizer. STEAM supports 133 candidate languages covering high-, medium-, and low-resource settings. Bayesian optimisation caps the number of evaluations at 20 per input, making the search tractable at this scale. The results show large and consistent performance gains over semantic clustering, with *average improvements of +0.23 AUC and +37.4 percentage points (%p) in TPR@1%*. We perform an extensive robustness analysis and adaptive adversarial evaluation, all confirming STEAM’s stability and effectiveness across diverse attack scenarios while maintaining a low false-positive

rate.

Our contributions are:


1. **Extensive multilingual evaluation.** We conduct a large-scale evaluation of multilingual watermarking methods across high-, medium-, and low-resource languages, uncovering weaknesses overlooked in prior work, which has focused exclusively on high-resource languages.
2. **Analysis of the limitations of semantic clustering.** We identify that the limitations of current multilingual watermarking stem from their core reliance on clusters of tokens.
3. **STEAM: a search-based, robust multilingual defence.** We introduce STEAM  $\hat{\mathcal{M}}$ , which applies Bayesian optimisation to select the back-translation language that best recovers watermark strength. STEAM already supports 133 candidate languages and is retroactively extensible to new ones, compatible with any watermarking technique and tokenizer, and non-invasive to the model output.
4. **Robustness across diverse languages.** STEAM  $\hat{\mathcal{M}}$  consistently outperforms existing multilingual watermarking methods, with improvements of up to **+0.41 AUC** and **+58.8% p TPR@1%**.

## 2. Related Work

Depending on when the watermark is applied, LLM watermarking techniques are generally classified into training-time watermarking and inference-time watermarking (also known as logit-based watermarking) (Liu et al., 2024b). This work focuses exclusively on the latter.

**Logit-based watermarking.** Logit-based watermarking embeds a watermark by directly modifying the token probability distribution (logits) during text generation (Liu et al., 2024b). The seminal approach, KGW (Kirchenbauer et al., 2023), partitions the tokenizer vocabulary into green and red lists using a random seed derived from a fixed window of previous tokens and biases generation towards green tokens. Zhao et al. (2023) proposed Unigram Watermarking, an extension of KGW that employs a fixed green/red partition to improve robustness against text editing and paraphrasing attacks. To maintain text quality, Hu et al. (2023) introduced an unbiased watermarking approach that integrates watermarks without altering the overall probability distribution of the output. Several works (Lee et al., 2024; Lu et al., 2024a; Liu & Bu, 2024; Wu et al., 2024) further improve robustness while preserving text quality.

Beyond these, ITS and EXP (Kuditipudi et al., 2024) offer model-agnostic, distortion-free watermarking schemes that

remain robust to text manipulation attacks. Our work analyses the multilingual capabilities of these techniques and builds upon them to develop our defence, STEAM .

**Watermarking robustness.** Several studies, including SIR (Liu et al., 2024a), SemaMark (Ren et al., 2024), semantic-aware watermarking (Fu et al., 2024), and SempStamp (Hou et al., 2024), incorporate semantic information to improve the robustness of watermarks against text transformation attacks. To achieve a balanced and context-aware partitioning of the green and red token lists, Guo et al. (2024) leveraged locality-sensitive hashing (LSH) (Indyk & Motwani, 1998) to generate a semantic key from contextual embeddings. Inspired by the inherent redundancy of multimedia data, WatME (Chen et al., 2024) embeds mutual exclusion rules within the lexical space for text watermarking. Furthermore, Luo et al. (2025) identified watermark collision, where multiple watermarks interact in ways that distort statistical distributions and hinder detection. These approaches primarily target attacks within a single language, such as paraphrasing or text editing, where the overall token distribution remains broadly stable. They do not address translation attacks, which shift text to a different language and cause a substantially different token distribution. This is a distinct failure mode that requires a different defence.

**Multilingual watermarking.** While much of the initial research focused on monolingual English text, a growing body of work now addresses the unique challenges of cross-lingual watermarking. A foundational contribution in this area is X-SIR (He et al., 2024), a direct extension of the SIR framework designed to defend against translation attacks. Other works have focused on evaluating the cross-lingual robustness of existing methods. For example, Han et al. (2025) assessed the robustness of SynthID-Text (Dathathri et al., 2024) to meaning-preserving transformations like back-translation. Similarly, Al Ghanim et al. (2025) conduct a comparative evaluation of four watermarking methods: KGW, Unigram, EXP, and X-SIR. Their analysis assesses robustness and text quality under various parameters and removal attacks in cross-lingual settings. Although these studies provide valuable insights, their scope is often limited to high-resource languages. Our work addresses this gap by providing a more comprehensive cross-lingual evaluation that includes an extensive set of low- and medium-resource languages.

### 3. Experimental Setup

This section outlines the experimental setup used to assess the robustness of multilingual watermarking methods across different languages, models, and attack scenarios.


Criterion	KGW & SIR	X-KGW & X-SIR	STEAM  (ours)
<i>Multilingual support</i>	✗	✓	✓
Non-invasive	–	✗	✓
Watermark-agnostic	–	✗	✓
Tokenizer-agnostic	–	✗	✓
<i>New language support</i>			
Medium-resource	–	~	✓
Low-resource	–	✗	✓ <sup>†</sup>
Retroactive support	–	✗	✓

Table 1. Comparison of watermarking methods and their multilingual capabilities. Criteria definitions in Section A.5.

✓ = Yes, ✗ = No, ~ = Limited, – = Not applicable  
<sup>†</sup> Requires translator (low-quality translation sufficient)

**Dataset.** We base our evaluation on the English subset of the mC4 dataset (Raffel et al., 2023), following the setup introduced by He et al. (2024). We sample a test set of 500 texts for all experiments.

**Attacks.** We evaluate watermark robustness under two translation-based attacks. Unless specified otherwise, all translations are performed with Google Translate. The first, direct translation, converts English outputs into a target language and is used in the main experiments. The second applies multi-step translation through a pivot language (He et al., 2024) and is reported in Appendix B.

**Multilingual models.** We use the following multilingual language models: `Aya-23-8B` (Aryabumi et al., 2024), `LLaMA-3.2-1B` (Grattafiori et al., 2024), and `LLaMAX-8B` (Lu et al., 2024b).

**Watermarking methods.** We analyse three watermarking schemes. We use the standard KGW (Kirchenbauer et al., 2023) as our primary non-multilingual baseline. Second, we evaluate X-SIR (He et al., 2024), a foundational work that proposes semantic clustering for cross-lingual robustness. Finally, we introduce X-KGW, a method that applies semantic clustering to KGW. This setup allows us to isolate and measure the precise impact of semantic clustering on watermark robustness (see Appendix A.3 for details about X-KGW).

**Evaluation metrics.** We assess the strength of the watermark using two standard binary classification metrics: (i) Area Under the ROC Curve (*AUC*), measuring the probability that a watermarked sample receives a higher detection score than a non-watermarked one; and (ii) True Positive Rate at a fixed False Positive Rate (*TPR@1%*), the proportion of correctly identified watermarked texts when the false positive rate is fixed at 1%.

## 4. Semantic Clustering Fails in Diverse Multilingual Settings

We show that semantic clustering is not inherently multilingual: it lacks robustness in unsupported languages, and extending its coverage to more languages fails in medium- and low-resource settings. We then identify the structural cause of these failures.

### 4.1. Robustness Against Unsupported Languages

Semantic clustering has only been evaluated on the languages it explicitly supports, so its robustness in unsupported languages remains unknown. We assess semantic clustering both within its originally supported languages using a hold-out setting, and on a broader set of unsupported ones to evaluate its cross-lingual generalisation.

**Hold-one-out setup.** This experiment evaluates how strongly X-SIR depends on its set of supported languages to be robust. Using the same languages as He et al. (2024), we exclude one language from the semantic clustering and then test the method on that withheld language. This setup allows us to measure how much X-SIR’s robustness depends on explicit language support. The full results for all languages and models are provided in Appendix B.1. We find that excluding a language from the supported set leads to only minor average changes in performance: AUC decreases by  $-0.025$  and  $\text{TPR}@1\%$  by  $-0.036$  for LLaMA-3.2 1B, and by  $+0.009$  and  $-0.015$  respectively for Aya-23 8B. In several cases, AUC even increases when a language is removed (10 out of 16 for LLaMA-3.2 and 7 out of 16 for Aya), revealing that X-SIR’s behaviour is highly variable and its robustness unreliable.

**New languages setup.** This experiment evaluates how much X-SIR relies on explicit language support to remain robust. If there is a large enough overlap of words between languages, supporting all languages may not be necessary. To test this, we extend the evaluation to the following set of unsupported languages: Italian (it), Spanish (es), Portuguese (pt), Polish (pl), Dutch (nl), Croatian (hr), Czech (cs), Danish (da), Korean (ko), and Arabic (ar). Appendix B.2 reports the performance of X-SIR and X-KGW for Aya-23 and the other models. Overall performance remains relatively low for X-SIR, with average AUC and  $\text{TPR}@1\%$  of 0.75 and 0.14 for Aya-23 8B, and 0.675 and 0.07 for LLaMA-3.2 1B. Similar trends are observed for X-KGW. More importantly, several languages show clearly weaker watermark strength: for X-SIR, Arabic is the most vulnerable for Aya-23 8B (AUC of 0.687,  $\text{TPR}@1\%$  of 0.093), while Portuguese and Arabic are weakest for LLaMA-3.2 1B (AUC of 0.650,  $\text{TPR}@1\%$  of 0.055). These results indicate that even a single poorly supported language can allow an attacker to

bypass watermark detection, highlighting the fragility of semantic-clustering-based multilingual watermarking.

*Table 2. Semantic clustering is weak for unsupported languages.* Watermark strength (AUC and  $\text{TPR}@1\%$ ) of X-SIR and X-KGW, limited to the five originally supported languages (en, fr, de, zh, ja). Aya-23 8B generates English text, which is then translated into a new unsupported language for evaluation. Minimum marks the weakest robustness (best attack case).

New Lang.	X-SIR ( $\uparrow$ )		X-KGW ( $\uparrow$ )	
	AUC	$\text{TPR}@1\%$	AUC	$\text{TPR}@1\%$
it	0.796	0.177	0.772	0.238
es	0.754	0.155	0.807	0.230
pt	0.775	0.133	0.792	0.286
pl	0.749	0.127	0.762	0.236
nl	0.776	0.164	0.808	0.314
hr	0.726	0.124	0.757	0.210
cs	0.773	0.111	0.754	0.254
da	0.734	0.161	0.764	0.266
ko	0.729	0.136	0.754	0.226
ar	0.687	0.093	0.765	0.168
Min.	0.687 (ar)	0.093 (ar)	0.754 (cs, ko)	0.168 (ar)

### 4.2. Failure to Support a Broad Range of Languages

Since X-SIR and X-KGW are not robust against translation to some unsupported languages, one possible solution is to extend the set of supported languages to cover most languages. In this section, we show that even when more languages are explicitly included, neither method achieves consistent robustness.

To evaluate the effectiveness of semantic clustering across languages, we extend the support of X-SIR and X-KGW to 17 languages spanning high-, medium-, and low-resource settings (methodology in Appendix A.4). The high-resource group includes French, German, Italian, Spanish, and Portuguese; the medium-resource group includes Polish, Dutch, Russian, Hindi, Korean, and Japanese; and the low-resource group includes Bengali, Persian, Vietnamese, Hebrew, Ukrainian, and Tamil.

The results for Aya-23 8B are reported in Table 3. For high-resource languages, X-SIR reaches an average AUC of 0.786 and  $\text{TPR}@1\%$  of 0.158, while X-KGW achieves 0.792 and 0.311, respectively. These scores drop for medium-resource languages to 0.721 and 0.126 for X-SIR, and 0.753 and 0.260 for X-KGW. The decline continues for low-resource languages, where X-SIR records an average AUC of 0.689 and  $\text{TPR}@1\%$  of 0.109, and X-KGW reaches 0.739 and 0.224. This trend indicates that semantic clustering robustness depends on language resource availability.

The performance gap becomes even more pronounced for specific low-resource languages. Tamil (ta) represents the weakest case of X-SIR on Aya-23 with an AUC of 0.560 and  $\text{TPR}@1\%$  of 0.049. LLaMAX-3 shares the same ob-

Table 3. Even when more languages are explicitly supported, the robustness of semantic clustering decreases from high- to low-resource languages. We extend semantic clustering to 17 newly supported languages. Aya-23 8B generates a text in English, then the translation attack is applied using each of these supported languages. Minimum indicates the worst-case robustness, i.e., the best language for an attack. Other models in Appendix B.3.

Translation Attack		X-SIR (↑)		X-KGW (↑)	
Type	Lang.	AUC	TPR@1%	AUC	TPR@1%
High-resource	fr	0.791	0.149	0.787	0.280
	de	0.784	0.163	0.811	0.312
	it	0.798	0.152	0.789	0.354
	es	0.780	0.150	0.794	0.278
	pt	0.779	0.176	0.778	0.330
Medium-resource	pl	0.752	0.146	0.767	0.312
	nl	0.823	0.213	0.842	0.332
	ru	0.738	0.122	0.711	0.246
	hi	0.616	0.056	0.739	0.194
	ko	0.719	0.115	0.770	0.318
ja	0.679	0.103	0.688	0.160	
Low-resource	bn	0.622	0.055	0.711	0.180
	fa	0.726	0.131	0.734	0.242
	vi	0.762	0.157	0.778	0.308
	iw	0.725	0.115	0.745	0.220
	uk	0.738	0.148	0.731	0.222
	ta	0.560	0.049	0.737	0.172
Minimum		0.560 (ta)	0.049 (ta)	0.688 (ja)	0.160 (ja)

ervation (AUC of 0.561, TPR@1% of 0.067). Such drastic degradation highlights that even with explicit support, X-SIR and X-KGW fail to maintain reliable watermark detection across all languages.

These findings raise a critical question: why does explicit language support fail to guarantee robustness for semantic clustering-based watermarking?

### 4.3. On the Fundamental Limitations of Semantic Clustering in Multilingual Watermarking

Both X-SIR and X-KGW show clear weaknesses in mid- and low-resource languages. In this section, we argue that these limitations stem from a fundamental property of semantic clustering: its inability to generalise across languages due to the uneven coverage of full-word tokens in tokenizers.

Semantic clustering assigns watermark signals using multilingual dictionaries to group semantically equivalent words across languages (He et al., 2024). However, the share of dictionary words that appear as full tokens in tokenizer vocabularies varies sharply across languages (Appendix ??). Low-resource languages have very few full-word tokens, as low as 0.13% with Hebrew. BPE-based tokenizers allocate tokens by frequency in the training data, inherently favour-




Figure 2. Tokenizer vocabulary favours high-resource languages. Percentage of words in multilingual dictionaries that appear in the tokenizer vocabulary.




Figure 3. Languages with larger tokenizer vocabularies have higher watermark robustness. Average AUC per language and model across three seeds. Lines are least squared regressions.

ing high-resource languages and fragmenting others into subword units with limited semantic meaning.

Figure 3 shows the relationship between watermark robustness (AUC) and the number of full-word tokens in the tokenizer vocabulary for each language. Across all three models, we observe a clear positive correlation: languages with higher token coverage achieve stronger watermark robustness, while those with lower coverage are far more vulnerable. This reveals a fundamental limitation of semantic clustering: (i) In the extreme case where a language has no full-word tokens, X-KGW collapses to KGW, as no token clusters can be formed. (ii) Even multilingual tokenizers cannot fully resolve this issue, since BPE allocation inherently disadvantages underrepresented languages. (iii) Most importantly, this vulnerability extends to monolingual watermarking: when text generated in one language (e.g., English) is translated into another, the target language may

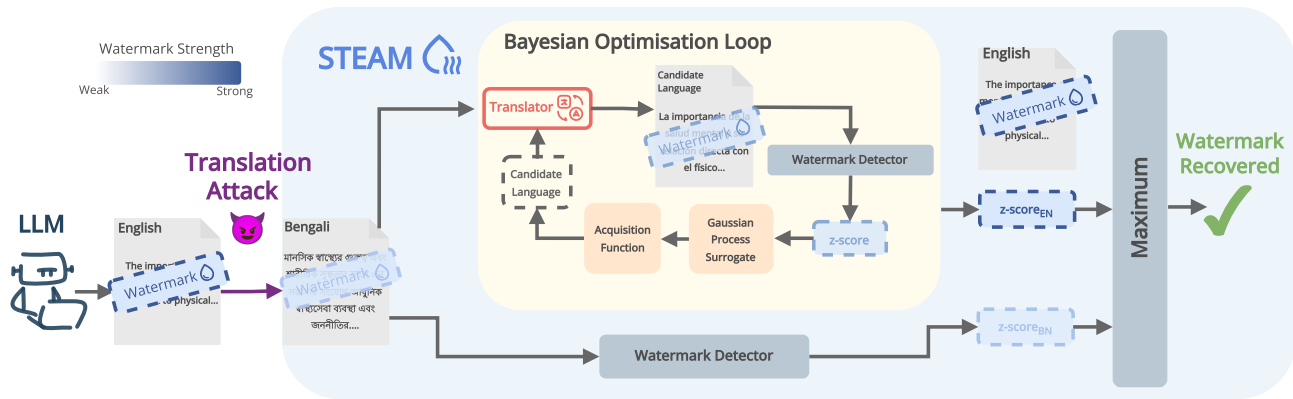


Figure 4. Overview of STEAM. A suspect text (Bengali) is back-translated into candidate languages selected by Bayesian optimisation from a pool of 133 languages. Each language is represented by continuous syntactic and phonological features, which guide the search. A standard watermark detector computes a  $z$ -score for each candidate text. STEAM selects the back-translation language that yields the highest  $z$ -score, here English, the language of the watermarked text before the translation attack.

contain far fewer full-word tokens, enabling the watermark to be lost. These findings underscore that the shortcomings of semantic clustering are structural and cannot be overcome by simply expanding language support or retraining tokenizers.

## 5. STEAM: Scaling Defence to 100+ Languages via Back-Translation

To address the limitations of semantic clustering for multilingual watermarking, we propose **Search-based Translation-Enhanced Approach for Multilingual watermarking (STEAM)**, a novel, model-agnostic defence method that uses Bayesian optimisation to search for the back-translation language that best recovers watermark strength. We first introduce STEAM, then evaluate its effectiveness across different adversarial scenarios, and finally analyse its robustness.

### 5.1. STEAM Description

STEAM recovers a watermark strength degraded by translation attacks via multilingual back-translation. For each suspect text, STEAM searches for the candidate language that, when used to back-translate the text, best recovers the watermark strength. Each candidate is evaluated using a watermark detector with a language-specific null hypothesis (described below). Finally, the highest corrected  $z$ -score across all evaluated candidates serves as the final test statistic. Figure 4 provides an overview of the STEAM pipeline.

**Bayesian optimisation for back-translation language selection.** We use Bayesian Optimisation (BO) to search for the back-translation language that best recovers watermark strength. Each language is characterised by a 131-dimensional feature vector with syntactic and phonolog-

ical properties sourced from URIEL (Khan et al., 2025), a knowledge base of linguistic properties. BO first evaluates a small set of randomly selected candidates, then fits a Gaussian process surrogate that models the relationship between linguistic features and observed  $z$ -scores, allowing it to predict which unevaluated languages are likely to yield high  $z$ -scores based on their similarity to already evaluated ones. At each subsequent iteration, the next back-translation language is chosen by maximising the expected improvement over all unevaluated candidates. The process repeats until a fixed budget of 20 evaluations is exhausted, enabling STEAM to scale to 133 candidate languages at low cost.

**Language-specific null hypothesis.** The standard  $z$ -score formula (Kirchenbauer et al., 2023) tests against a null hypothesis of a fixed green token fraction  $\gamma$  (e.g.  $\gamma = 50\%$ ), assumed uniform across all languages. This assumption is violated in low-resource languages, where tokenizers fragment single UTF-8 characters into sub-character tokens that are reused across many characters, making them disproportionately frequent. For instance, a single token accounts for 21.5% of all tokens in our Tamil texts under Llama 3 tokenizer (see Appendix B.7). If such a frequent token falls in the green list, the green token fraction exceeds  $\gamma$  and the  $z$ -score inflates; when it falls in the red list, the  $z$ -score deflates. In both cases, the shift is independent of any watermark strength. Since STEAM optimises for the highest  $z$ -score across candidate languages, such bias would cause it to systematically favour certain languages regardless of the watermark presence.

We address language token bias by replacing  $\gamma$  with a language- and key-specific  $\gamma_\ell$ . Concretely,  $\gamma_\ell$  is the empirical green token fraction measured on a calibration set of 500 human-written texts per candidate language, for a fixed watermark key:

Table 4. STEAM  $\hat{C}_m$  is consistently better than semantic clustering by a large margin. Watermark strength (AUC and TPR@1%) of multilingual watermarking techniques with 17 supported languages and Aya-23 8B. Red indicates robustness lower than the KGW baseline. Bolded is best. Other models in Appendix B.4.

Translation Attack		AUC ( $\uparrow$ )				TPR@1% ( $\uparrow$ )			
Type	Language	KGW	X-KGW	X-SIR	STEAM $\hat{C}_m$	KGW	X-KGW	X-SIR	STEAM $\hat{C}_m$
High-resource	fr	0.746	0.787	0.791	<b>0.976</b>	0.224	0.280	0.149	<b>0.582</b>
	de	0.730	0.811	0.784	<b>0.973</b>	0.224	0.312	0.163	<b>0.622</b>
	it	0.733	0.789	0.798	<b>0.978</b>	0.202	0.354	0.152	<b>0.530</b>
	es	0.717	0.794	0.780	<b>0.976</b>	0.232	0.278	0.150	<b>0.580</b>
	pt	0.733	0.778	0.779	<b>0.977</b>	0.246	0.330	0.176	<b>0.538</b>
Medium-resource	pl	0.729	0.767	0.752	<b>0.975</b>	0.228	0.312	0.146	<b>0.526</b>
	nl	0.767	0.842	0.823	<b>0.982</b>	0.290	0.332	0.213	<b>0.612</b>
	ru	0.667	0.711	0.738	<b>0.971</b>	0.158	0.246	0.122	<b>0.576</b>
	hi	0.614	0.739	0.616	<b>0.978</b>	0.120	0.194	0.056	<b>0.644</b>
	ko	0.730	0.770	0.719	<b>0.968</b>	0.210	0.318	0.115	<b>0.460</b>
Low-resource	ja	0.656	0.688	0.679	<b>0.975</b>	0.114	0.160	0.103	<b>0.498</b>
	bn	0.667	0.711	0.622	<b>0.978</b>	0.068	0.180	0.055	<b>0.604</b>
	fa	0.704	0.734	0.726	<b>0.979</b>	0.196	0.242	0.131	<b>0.664</b>
	vi	0.702	0.778	0.762	<b>0.976</b>	0.186	0.308	0.157	<b>0.626</b>
	iw	0.716	0.745	0.725	<b>0.974</b>	0.172	0.220	0.115	<b>0.578</b>
	uk	0.674	0.731	0.738	<b>0.979</b>	0.210	0.222	0.148	<b>0.542</b>
	ta	0.575	0.737	0.560	<b>0.967</b>	0.082	0.172	0.049	<b>0.504</b>

$$z_\ell = (n_g - \gamma_\ell n) / \sqrt{n\gamma_\ell(1 - \gamma_\ell)} \quad (1)$$

where  $n_g$  is the number of green tokens of the suspect text and  $n$  its total number of tokens.

## 5.2. STEAM Evaluation

**Comparison to semantic clustering.** We evaluate STEAM against semantic clustering methods to assess its robustness under translation attacks. As in §4.2, all methods are tested on the same set of 17 attack languages. For back-translation, STEAM chooses from a pool of 133 candidate languages covering a wide range of language families (Appendix A.7). STEAM achieves consistently strong results, maintaining an average AUC above **0.965** across all language categories, including medium- and low-resource ones (Table 4, Appendix B.4). Compared with semantic clustering approaches (X-SIR and X-KGW), STEAM shows large gains: on average, +0.25 AUC and +44.0%p TPR@1% relative to X-SIR, and +0.216 AUC and +30.7%p TPR@1% relative to X-KGW. The largest improvements are observed for Tamil and Hindi, with up to +0.41 AUC and +58.8%p TPR@1%, respectively. These gains confirm that STEAM generalises reliably beyond high-resource settings. Unlike semantic clustering, STEAM is robust in medium- and low-resource languages, unaffected by tokenizer limitations.

**STEAM robustness to unsupported languages.** We next evaluate whether STEAM remains robust when the best language is not in its pool of candidate languages. This is a realistic scenario, since an attacker may use any language.

Table 5. STEAM  $\hat{C}_m$  remains robust on unsupported languages and outperforms multilingual methods by a large margin. Bolded is best. Red indicates lower than the KGW baseline. Full table in Appendix B.4.

New Language	AUC ( $\uparrow$ )			
	KGW	X-KGW	X-SIR	STEAM $\hat{C}_m$
it	0.733	0.772	0.796	<b>0.966</b>
es	0.717	0.807	0.754	<b>0.967</b>
pt	0.732	0.792	0.775	<b>0.971</b>
pl	0.730	0.762	0.749	<b>0.960</b>
nl	0.768	0.808	0.776	<b>0.966</b>
hr	0.706	0.757	0.726	<b>0.965</b>
cs	0.717	0.754	0.773	<b>0.974</b>
da	0.713	0.764	0.734	<b>0.971</b>
ko	0.732	0.754	0.729	<b>0.961</b>
ar	0.689	0.765	0.687	<b>0.971</b>

To ensure a fair comparison, we remove the language of the watermarked text (before the translation attack) from all methods: for semantic clustering, we exclude the corresponding dictionary; for STEAM, we remove the language from the back-translation pool. In this setup, STEAM maintains strong detection across all evaluated languages, with an average AUC of 0.967 (Table 5). It outperforms semantic clustering by a large margin, with average gains of +0.19 AUC and +20.6%p TPR@1% over X-KGW and +0.22 AUC and +31.0%p TPR@1% over X-SIR. Unlike semantic clustering methods, which lose coverage when the language is absent from their dictionaries, STEAM remains resilient: linguistically similar languages in its back-translation pool are sufficient to recover the watermark strength.

Table 6. STEAM  $\mathbb{C}_m$  is robust to translator mismatch. Average AUC across German, Hindi, and Hebrew for all attacker and defender translator combinations. *Attacker* translator is used for translation attack. *Defender* translator is used by STEAM. DS-V3.2 stands for DeepSeek-V3.2-Exp.

Attacker	Defender		
	G. Translate	DS-V3.2	GPT-4o-mini
Google Translate	0.976	0.974	0.975
DS-V3.2	0.979	0.952	0.943
GPT-4o-mini	0.976	0.973	0.964

### 5.3. Robustness Analysis

**Robustness to translator mismatch.** The robustness of STEAM should not depend on the specific translation service used. An adversary could try to bypass our defence by using a different translation system for their attack. We examine whether the performance of STEAM remains robust when the attacker and defender use different translators. Table 6 reports the average AUC across German, Hindi, and Hebrew for all nine combinations of three translation services (Google Translate, DeepSeek-V3.2-Exp (DeepSeek-AI et al., 2025), and GPT-4o-mini). All pairs achieve an AUC above 0.94, and the diagonal (matched translator) does not consistently outperform mismatched settings. This suggests that our method genuinely recovers watermark strength independently of the specific translation service rather than relying on translator-specific artefacts.

**Adaptive evaluation: multistep translation attack.** To assess the robustness of our defence under adaptive attack, we introduce a stronger multistep translation attack that adds an extra translation step beyond the single-hop setup. This design prevents STEAM from relying on direct back-translation to recover the watermark strength. In this two-step attack, the text is first translated using the full set of languages from §4.2, and the resulting output is then translated again through one of three pivot languages: German (high-resource), Korean (medium-resource), or Bengali (low-resource) (Appendix B.4). Despite this adaptive setup, STEAM remains robust, maintaining an average AUC of 0.884 across all conditions. The lowest result occurs with a Korean pivot on high-resource languages (AUC of 0.833), while all other settings remain above 0.860. Although such multi-hop attacks can weaken other defences, they also tend to reduce overall translation quality, limiting their practical impact.

## 6. Conclusion

We showed that current multilingual watermarking methods fail to remain robust under translation attacks, especially in medium- and low-resource languages. If watermarking

lacks robustness in a given language, online content in that language may be disproportionately affected by synthetic or undesirable content. This risk is especially serious for low- and medium-resource languages, which already face a shortage of high-quality digital resources and often lack effective moderation systems.

To address this, we introduced STEAM  $\mathbb{C}_m$ , a watermark-agnostic method that applies Bayesian optimisation to select the back-translation language that best recovers watermark strength. Extensive experiments scaling to 133 languages and diverse attack scenarios show that STEAM achieves consistently stronger robustness and fairness than existing multilingual watermarking methods, particularly in medium- and low-resource settings.

Our findings highlight the need for watermarking research to treat linguistic diversity as a core requirement, ensuring that the security and trust of large language models extend to all languages, not only those with abundant digital resources.

## 440 Limitations

441 While our proposed method, STEAM, demonstrates significant improvements in multilingual watermarking, we acknowledge several limitations that also present avenues for future research.

442 Our evaluation considers a set of 17 languages out of the 133 supported, chosen to represent diverse linguistic families. However, this set might not be fully representative of the linguistic diversity of the world.

443 STEAM demonstrates clear advantages over prior multilingual watermarking techniques on supported languages. However, its performance on unsupported languages remains comparable to existing methods. Nevertheless, a key strength of STEAM is that it can easily support additional languages. We believe that a broad coverage of languages is necessary for all multilingual watermark techniques.

444 The operational cost of STEAM, measured in translation API requests, is bounded by the Bayesian optimisation budget of 20 evaluations per input, regardless of the number of supported languages. Our empirical results show that the method’s performance gains do not depend on high-cost translation services. The use of standard, widely available tools like Google Translate proved sufficient to achieve consistent improvements.

445 Finally, the current implementation of STEAM is specifically designed to defend against translation-based attacks. It is not designed to counter other significant text transformation attacks, such as paraphrasing attacks. This focus is a deliberate choice: STEAM is designed to be modular, allowing the translation robustness component to operate independently. This modularity ensures that other parts of the watermarking pipeline are not affected and provides a clear path for future enhancements. Future research could focus on creating and integrating new modules to build a more holistically robust watermarking system.

## 446 Ethical Considerations

447 This work has potential dual-use implications. On one hand, studying adversarial attacks against watermarking could inform malicious actors about possible strategies to weaken watermark defences. However, we believe the benefits outweigh these risks.

448 First, our contribution is not only an analysis but also a concrete defence (STEAM) that achieves a high level of robustness, substantially exceeding prior multilingual watermarking methods. Our results demonstrate that STEAM provides consistently strong robustness against translation attacks across a wide range of languages.

449 Second, by explicitly addressing low- and medium-resource

languages, our method promotes fairness: watermarking becomes more reliable across diverse linguistic settings, rather than being limited to a handful of high-resource languages.

Robust multilingual watermarking is an important safeguard against misuse of large language models, such as the generation and dissemination of fake news or disinformation in less-resourced languages where moderation tools are often weaker. We view this work as a step toward improving the security and trustworthiness of multilingual AI systems.

## References

- Al Ghanim, M., Xue, J., Hastuti, R. P., Zheng, M., Solihin, Y., and Lou, Q. Evaluating the robustness and accuracy of text watermarking under real-world cross-lingual manipulations. In *Findings of the Association for Computational Linguistics: EMNLP 2025*. Association for Computational Linguistics, 2025. URL <https://arxiv.org/abs/2502.16699>.
- Aryabumi, V., Dang, J., Talupuru, D., Dash, S., Cairuz, D., Lin, H., Venkitesh, B., Smith, M., Marchisio, K., Ruder, S., Locatelli, A., Kreutzer, J., Frosst, N., Blunsom, P., Fadaee, M., Üstün, A., and Hooker, S. Aya 23: Open weight releases to further multilingual progress, 2024.
- Balandat, M., Karrer, B., Jiang, D. R., Daulton, S., Letham, B., Wilson, A. G., and Bakshy, E. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/f5b1b89d98b7286673128a5fb112cb9a-Abstract.html>.
- Chen, L., Bian, Y., Deng, Y., Cai, D., Li, S., Zhao, P., and fai Wong, K. Watme: Towards lossless watermarking through lexical redundancy, 2024. URL <https://arxiv.org/abs/2311.09832>.
- Chen, R., Wu, Y., Guo, J., and Huang, H. Improved unbiased watermark for large language models. In Che, W., Nabende, J., Shutova, E., and Pilehvar, M. T. (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 20587–20601, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.1005. URL <https://aclanthology.org/2025.acl-long.1005/>.
- Conneau, A., Lample, G., Ranzato, M., Denoyer, L., and Jégou, H. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*, 2017.

- 495 Dathathri, S., See, A., Ghaisas, S., Huang, P.-S., McAdam,  
496 R., Welbl, J., Bachani, V., Kaskasoli, A., Stan-  
497 forth, R., Matejovicova, T., Hayes, J., Vyas, N.,  
498 Merey, M. A., Brown-Cohen, J., Bunel, R., Balle,  
499 B., Cemgil, T., Ahmed, Z., Stacpoole, K., Shumailov,  
500 I., Baetu, C., Goyal, S., Hassabis, D., and Kohli,  
501 P. Scalable watermarking for identifying large lan-  
502 guage model outputs. *Nature*, 634(8035):818–823,  
503 October 2024. ISSN 1476-4687. doi: 10.1038/  
504 s41586-024-08025-4. URL [http://dx.doi.org/  
505 10.1038/s41586-024-08025-4](http://dx.doi.org/10.1038/s41586-024-08025-4).  
506
- 507 DeepSeek-AI, Liu, A., Feng, B., Xue, B., Wang, B., Wu, B.,  
508 Lu, C., Zhao, C., Deng, C., Zhang, C., Ruan, C., Dai, D.,  
509 Guo, D., Yang, D., Chen, D., Ji, D., Li, E., Lin, F., Dai,  
510 F., Luo, F., Hao, G., Chen, G., Li, G., Zhang, H., Bao,  
511 H., Xu, H., Wang, H., Zhang, H., Ding, H., Xin, H., Gao,  
512 H., Li, H., Qu, H., Cai, J. L., Liang, J., Guo, J., Ni, J., Li,  
513 J., Wang, J., Chen, J., Chen, J., Yuan, J., Qiu, J., Li, J.,  
514 Song, J., Dong, K., Hu, K., Gao, K., Guan, K., Huang,  
515 K., Yu, K., Wang, L., Zhang, L., Xu, L., Xia, L., Zhao,  
516 L., Wang, L., Zhang, L., Li, M., Wang, M., Zhang, M.,  
517 Zhang, M., Tang, M., Li, M., Tian, N., Huang, P., Wang,  
518 P., Zhang, P., Wang, Q., Zhu, Q., Chen, Q., Du, Q., Chen,  
519 R. J., Jin, R. L., Ge, R., Zhang, R., Pan, R., Wang, R.,  
520 Xu, R., Zhang, R., Chen, R., Li, S. S., Lu, S., Zhou, S.,  
521 Chen, S., Wu, S., Ye, S., Ye, S., Ma, S., Wang, S., Zhou,  
522 S., Yu, S., Zhou, S., Pan, S., Wang, T., Yun, T., Pei, T.,  
523 Sun, T., Xiao, W. L., Zeng, W., Zhao, W., An, W., Liu,  
524 W., Liang, W., Gao, W., Yu, W., Zhang, W., Li, X. Q.,  
525 Jin, X., Wang, X., Bi, X., Liu, X., Wang, X., Shen, X.,  
526 Chen, X., Zhang, X., Chen, X., Nie, X., Sun, X., Wang,  
527 X., Cheng, X., Liu, X., Xie, X., Liu, X., Yu, X., Song,  
528 X., Shan, X., Zhou, X., Yang, X., Li, X., Su, X., Lin, X.,  
529 Li, Y. K., Wang, Y. Q., Wei, Y. X., Zhu, Y. X., Zhang,  
530 Y., Xu, Y., Xu, Y., Huang, Y., Li, Y., Zhao, Y., Sun, Y.,  
531 Li, Y., Wang, Y., Yu, Y., Zheng, Y., Zhang, Y., Shi, Y.,  
532 Xiong, Y., He, Y., Tang, Y., Piao, Y., Wang, Y., Tan, Y.,  
533 Ma, Y., Liu, Y., Guo, Y., Wu, Y., Ou, Y., Zhu, Y., Wang,  
534 Y., Gong, Y., Zou, Y., He, Y., Zha, Y., Xiong, Y., Ma, Y.,  
535 Yan, Y., Luo, Y., You, Y., Liu, Y., Zhou, Y., Wu, Z. F.,  
536 Ren, Z. Z., Ren, Z., Sha, Z., Fu, Z., Xu, Z., Huang, Z.,  
537 Zhang, Z., Xie, Z., Zhang, Z., Hao, Z., Gou, Z., Ma, Z.,  
538 Yan, Z., Shao, Z., Xu, Z., Wu, Z., Zhang, Z., Li, Z., Gu,  
539 Z., Zhu, Z., Liu, Z., Li, Z., Xie, Z., Song, Z., Gao, Z.,  
540 and Pan, Z. Deepseek-v3 technical report, 2025. URL  
541 <https://arxiv.org/abs/2412.19437>.  
542
- 543 Fu, Y., Xiong, D., and Dong, Y. Watermarking conditional  
544 text generation for ai detection: Unveiling challenges  
545 and a semantic-aware watermark remedy, 2024. URL  
546 <https://arxiv.org/abs/2307.13808>.  
547
- 548 Gardner, J. R., Pleiss, G., Bindel, D., Weinberger, K. Q.,  
549 and Wilson, A. G. Gpytorch: Blackbox matrix-matrix  
gaussian process inference with gpu acceleration, 2021.  
URL <https://arxiv.org/abs/1809.11165>.
- Grattafiori, A., Dubey, A., Jauhri, A., Pandey, A., Kadian,  
A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A.,  
Vaughan, A., Yang, A., Fan, A., Goyal, A., Hartshorn,  
A., Yang, A., Mitra, A., Sravankumar, A., Korenev,  
A., Hinsvark, A., Rao, A., Zhang, A., Rodriguez, A.,  
Gregerson, A., Spataru, A., Roziere, B., Biron, B., Tang,  
B., Chern, B., Caucheteux, C., Nayak, C., Bi, C., Marra,  
C., McConnell, C., Keller, C., Touret, C., Wu, C., Wong,  
C., Ferrer, C. C., Nikolaidis, C., Allonsius, D., Song, D.,  
Pintz, D., Livshits, D., Wyatt, D., Esiobu, D., Choudhary,  
D., Mahajan, D., Garcia-Olano, D., Perino, D., Hupkes,  
D., Lakomkin, E., AlBadawy, E., Lobanova, E., Dinan,  
E., Smith, E. M., Radenovic, F., Guzmán, F., Zhang, F.,  
Synnaeve, G., Lee, G., Anderson, G. L., Thattai, G., Nail,  
G., Mialon, G., Pang, G., Cucurell, G., Nguyen, H., Ko-  
revaar, H., Xu, H., Touvron, H., Zarov, I., Ibarra, I. A.,  
Kloumann, I., Misra, I., Evtimov, I., Zhang, J., Copet, J.,  
Lee, J., Geffert, J., Vranes, J., Park, J., Mahadeokar, J.,  
Shah, J., van der Linde, J., Billock, J., Hong, J., Lee, J.,  
Fu, J., Chi, J., Huang, J., Liu, J., Wang, J., Yu, J., Bitton,  
J., Spisak, J., Park, J., Rocca, J., Johnstun, J., Saxe, J., Jia,  
J., Alwala, K. V., Prasad, K., Upasani, K., Plawiak, K., Li,  
K., Heafield, K., Stone, K., El-Arini, K., Iyer, K., Malik,  
K., Chiu, K., Bhalla, K., Lakhotia, K., Rantala-Yearly,  
L., van der Maaten, L., Chen, L., Tan, L., Jenkins, L.,  
Martin, L., Madaan, L., Malo, L., Blecher, L., Landzaat,  
L., de Oliveira, L., Muzzi, M., Pasupuleti, M., Singh,  
M., Paluri, M., Kardas, M., Tsimpoukelli, M., Oldham,  
M., Rita, M., Pavlova, M., Kambadur, M., Lewis, M.,  
Si, M., Singh, M. K., Hassan, M., Goyal, N., Torabi, N.,  
Bashlykov, N., Bogoychev, N., Chatterji, N., Zhang, N.,  
Duchenne, O., Çelebi, O., Alrassy, P., Zhang, P., Li, P.,  
Vasic, P., Weng, P., Bhargava, P., Dubal, P., Krishnan,  
P., Koura, P. S., Xu, P., He, Q., Dong, Q., Srinivasan,  
R., Ganapathy, R., Calderer, R., Cabral, R. S., Stojnic,  
R., Raileanu, R., Maheswari, R., Girdhar, R., Patel, R.,  
Sauvestre, R., Polidoro, R., Sumbaly, R., Taylor, R., Silva,  
R., Hou, R., Wang, R., Hosseini, S., Chennabasappa, S.,  
Singh, S., Bell, S., Kim, S. S., Edunov, S., Nie, S., Narang,  
S., Raparthy, S., Shen, S., Wan, S., Bhosale, S., Zhang,  
S., Vandenhende, S., Batra, S., Whitman, S., Sootla, S.,  
Collot, S., Gururangan, S., Borodinsky, S., Herman, T.,  
Fowler, T., Sheasha, T., Georgiou, T., Scialom, T., Speck-  
bacher, T., Mihaylov, T., Xiao, T., Karn, U., Goswami, V.,  
Gupta, V., Ramanathan, V., Kerkez, V., Gonguet, V., Do,  
V., Vogeti, V., Albiero, V., Petrovic, V., Chu, W., Xiong,  
W., Fu, W., Meers, W., Martinet, X., Wang, X., Wang,  
X., Tan, X. E., Xia, X., Xie, X., Jia, X., Wang, X., Gold-  
schlag, Y., Gaur, Y., Babaei, Y., Wen, Y., Song, Y., Zhang,  
Y., Li, Y., Mao, Y., Coudert, Z. D., Yan, Z., Chen, Z.,  
Papakipos, Z., Singh, A., Srivastava, A., Jain, A., Kelsey,

- 550 A., Shajnfeld, A., Gangidi, A., Victoria, A., Goldstand,  
 551 A., Menon, A., Sharma, A., Boesenberg, A., Baevski, A.,  
 552 Feinstein, A., Kallet, A., Sangani, A., Teo, A., Yunus, A.,  
 553 Lupu, A., Alvarado, A., Caples, A., Gu, A., Ho, A., Poul-  
 554 ton, A., Ryan, A., Ramchandani, A., Dong, A., Franco,  
 555 A., Goyal, A., Saraf, A., Chowdhury, A., Gabriel, A.,  
 556 Bharambe, A., Eisenman, A., Yazdan, A., James, B.,  
 557 Maurer, B., Leonhardi, B., Huang, B., Loyd, B., Paola,  
 558 B. D., Paranjape, B., Liu, B., Wu, B., Ni, B., Hancock,  
 559 B., Wasti, B., Spence, B., Stojkovic, B., Gamido, B.,  
 560 Montalvo, B., Parker, C., Burton, C., Mejia, C., Liu, C.,  
 561 Wang, C., Kim, C., Zhou, C., Hu, C., Chu, C.-H., Cai, C.,  
 562 Tindal, C., Feichtenhofer, C., Gao, C., Civin, D., Beaty,  
 563 D., Kreymer, D., Li, D., Adkins, D., Xu, D., Testuggine,  
 564 D., David, D., Parikh, D., Liskovich, D., Foss, D., Wang,  
 565 D., Le, D., Holland, D., Dowling, E., Jamil, E., Mont-  
 566 gomery, E., Presani, E., Hahn, E., Wood, E., Le, E.-T.,  
 567 Brinkman, E., Arcaute, E., Dunbar, E., Smothers, E., Sun,  
 568 F., Kreuk, F., Tian, F., Kokkinos, F., Ozgenel, F., Cag-  
 569 gioni, F., Kanayet, F., Seide, F., Florez, G. M., Schwarz,  
 570 G., Badeer, G., Swee, G., Halpern, G., Herman, G., Sizov,  
 571 G., Guangyi, Zhang, Lakshminarayanan, G., Inan, H.,  
 572 Shojanazeri, H., Zou, H., Wang, H., Zha, H., Habeeb, H.,  
 573 Rudolph, H., Suk, H., Aspegren, H., Goldman, H., Zhan,  
 574 H., Damlaj, I., Molybog, I., Tufanov, I., Leontiadis, I.,  
 575 Veliche, I.-E., Gat, I., Weissman, J., Geboski, J., Kohli,  
 576 J., Lam, J., Asher, J., Gaya, J.-B., Marcus, J., Tang, J.,  
 577 Chan, J., Zhen, J., Reizenstein, J., Teboul, J., Zhong, J.,  
 578 Jin, J., Yang, J., Cummings, J., Carvill, J., Shepard, J.,  
 579 McPhee, J., Torres, J., Ginsburg, J., Wang, J., Wu, K., U,  
 580 K. H., Saxena, K., Khandelwal, K., Zand, K., Matosich,  
 581 K., Veeraraghavan, K., Michelena, K., Li, K., Jagadeesh,  
 582 K., Huang, K., Chawla, K., Huang, K., Chen, L., Garg,  
 583 L., A. L., Silva, L., Bell, L., Zhang, L., Guo, L., Yu, L.,  
 584 Moshkovich, L., Wehrstedt, L., Khabisa, M., Avalani, M.,  
 585 Bhatt, M., Mankus, M., Hasson, M., Lennie, M., Reso,  
 586 M., Groshev, M., Naumov, M., Lathi, M., Keneally, M.,  
 587 Liu, M., Seltzer, M. L., Valko, M., Restrepo, M., Patel,  
 588 M., Vyatskov, M., Samvelyan, M., Clark, M., Macey,  
 589 M., Wang, M., Hermoso, M. J., Metanat, M., Rastegari,  
 590 M., Bansal, M., Santhanam, N., Parks, N., White, N.,  
 591 Bawa, N., Singhal, N., Egebo, N., Usunier, N., Mehta,  
 592 N., Laptev, N. P., Dong, N., Cheng, N., Chernoguz, O.,  
 593 Hart, O., Salpekar, O., Kalinli, O., Kent, P., Parekh, P.,  
 594 Saab, P., Balaji, P., Rittner, P., Bontrager, P., Roux, P.,  
 595 Dollar, P., Zvyagina, P., Ratanchandani, P., Yuvraj, P.,  
 596 Liang, Q., Alao, R., Rodriguez, R., Ayub, R., Murthy, R.,  
 597 Nayani, R., Mitra, R., Parthasarathy, R., Li, R., Hogan,  
 598 R., Battey, R., Wang, R., Howes, R., Rinott, R., Mehta,  
 599 S., Siby, S., Bondu, S. J., Datta, S., Chugh, S., Hunt, S.,  
 600 Dhillon, S., Sidorov, S., Pan, S., Mahajan, S., Verma,  
 601 S., Yamamoto, S., Ramaswamy, S., Lindsay, S., Lindsay,  
 602 S., Feng, S., Lin, S., Zha, S. C., Patil, S., Shankar, S.,  
 603 Zhang, S., Zhang, S., Wang, S., Agarwal, S., Sajuyigbe,  
 604 S., Chintala, S., Max, S., Chen, S., Kehoe, S., Satter-  
 field, S., Govindaprasad, S., Gupta, S., Deng, S., Cho,  
 S., Virk, S., Subramanian, S., Choudhury, S., Goldman,  
 S., Remez, T., Glaser, T., Best, T., Koehler, T., Robinson,  
 T., Li, T., Zhang, T., Matthews, T., Chou, T., Shaked,  
 T., Vontimitta, V., Ajayi, V., Montanez, V., Mohan, V.,  
 Kumar, V. S., Mangla, V., Ionescu, V., Poenaru, V., Mi-  
 hailescu, V. T., Ivanov, V., Li, W., Wang, W., Jiang, W.,  
 Bouaziz, W., Constable, W., Tang, X., Wu, X., Wang, X.,  
 Wu, X., Gao, X., Kleinman, Y., Chen, Y., Hu, Y., Jia, Y.,  
 Qi, Y., Li, Y., Zhang, Y., Zhang, Y., Adi, Y., Nam, Y., Yu,  
 Wang, Zhao, Y., Hao, Y., Qian, Y., Li, Y., He, Y., Rait,  
 Z., DeVito, Z., Rosnbrick, Z., Wen, Z., Yang, Z., Zhao,  
 Z., and Ma, Z. The llama 3 herd of models, 2024. URL  
<https://arxiv.org/abs/2407.21783>.
- Guo, Y., Tian, Z., Song, Y., Liu, T., Ding, L., and Li,  
 D. Context-aware watermark with semantic balanced  
 green-red lists for large language models. In Al-Onaizan,  
 Y., Bansal, M., and Chen, Y.-N. (eds.), *Proceedings  
 of the 2024 Conference on Empirical Methods in Nat-  
 ural Language Processing*, pp. 22633–22646, Miami,  
 Florida, USA, November 2024. Association for Compu-  
 tational Linguistics. doi: 10.18653/v1/2024.emnlp-main.  
 1260. URL [https://aclanthology.org/2024.  
 emnlp-main.1260/](https://aclanthology.org/2024.<br/>
  emnlp-main.1260/).
- Han, X., Li, Q., Ni, J., and Zulkernine, M. Robustness  
 assessment and enhancement of text watermarking for  
 google’s synthid, 2025. URL [https://arxiv.org/  
 abs/2508.20228](https://arxiv.org/<br/>
  abs/2508.20228).
- He, Z., Zhou, B., Hao, H., Liu, A., Wang, X., Tu, Z., Zhang,  
 Z., and Wang, R. Can watermarks survive translation? on  
 the cross-lingual consistency of text watermark for large  
 language models. In Ku, L.-W., Martins, A., and Srikur-  
 mar, V. (eds.), *Proceedings of the 62nd Annual Meeting  
 of the Association for Computational Linguistics (Volume  
 1: Long Papers)*, pp. 4115–4129, Bangkok, Thailand,  
 August 2024. Association for Computational Linguis-  
 tics. doi: 10.18653/v1/2024.acl-long.226. URL [https://  
 aclanthology.org/2024.acl-long.226/](https://<br/>
  aclanthology.org/2024.acl-long.226/).
- Hou, A., Zhang, J., He, T., Wang, Y., Chuang, Y.-S., Wang,  
 H., Shen, L., Van Durme, B., Khashabi, D., and Tsvetkov,  
 Y. SemStamp: A semantic watermark with paraphras-  
 tic robustness for text generation. In Duh, K., Gomez,  
 H., and Bethard, S. (eds.), *Proceedings of the 2024 Con-  
 ference of the North American Chapter of the Associa-  
 tion for Computational Linguistics: Human Language  
 Technologies (Volume 1: Long Papers)*, pp. 4067–4082,  
 Mexico City, Mexico, June 2024. Association for Compu-  
 tational Linguistics. doi: 10.18653/v1/2024.naacl-long.  
 226. URL [https://aclanthology.org/2024.  
 naacl-long.226/](https://aclanthology.org/2024.<br/>
  naacl-long.226/).

- 605 Hu, Z., Chen, L., Wu, X., Wu, Y., Zhang, H., and Huang,  
606 H. Unbiased watermark for large language models, 2023.  
607 URL <https://arxiv.org/abs/2310.10669>.
- 608 Indyk, P. and Motwani, R. Approximate nearest neigh-  
609 bors: towards removing the curse of dimensionality. In  
610 *Proceedings of the Thirtieth Annual ACM Symposium on*  
611 *Theory of Computing*, STOC '98, pp. 604–613, New York,  
612 NY, USA, 1998. Association for Computing Machinery.  
613 ISBN 0897919629. doi: 10.1145/276698.276876. URL  
614 <https://doi.org/10.1145/276698.276876>.
- 615 Khan, A., Shipton, M., Anugraha, D., Duan, K., Hoang,  
616 P. H., Khiu, E., Dođruöz, A. S., and Lee, E.-S. A. Uriel+:  
617 Enhancing linguistic inclusion and usability in a typo-  
618 logical and multilingual knowledge base, 2025. URL  
619 <https://arxiv.org/abs/2409.18472>.
- 620 Kirchenbauer, J., Geiping, J., Wen, Y., Katz, J., Miers,  
621 I., and Goldstein, T. A watermark for large language  
622 models. In Krause, A., Brunskill, E., Cho, K., Engel-  
623 hardt, B., Sabato, S., and Scarlett, J. (eds.), *Pro-*  
624 *ceedings of the 40th International Conference on Ma-*  
625 *chine Learning*, volume 202 of *Proceedings of Machine*  
626 *Learning Research*, pp. 17061–17084. PMLR, 23–29 Jul  
627 2023. URL <https://proceedings.mlr.press/v202/kirchenbauer23a.html>.
- 628 Kudipudi, R., Thickstun, J., Hashimoto, T., and Liang,  
629 P. Robust distortion-free watermarks for language mod-  
630 els, 2024. URL <https://arxiv.org/abs/2307.15593>.
- 631 Lee, T., Hong, S., Ahn, J., Hong, I., Lee, H., Yun, S., Shin,  
632 J., and Kim, G. Who wrote this code? watermarking for  
633 code generation, 2024. URL <https://arxiv.org/abs/2305.15060>.
- 634 Littell, P., Mortensen, D. R., Lin, K., Kairis, K., Turner,  
635 C., and Levin, L. URIEL and lang2vec: Representing  
636 languages as typological, geographical, and phylogenetic  
637 vectors. In Lapata, M., Blunsom, P., and Koller, A. (eds.),  
638 *Proceedings of the 15th Conference of the European*  
639 *Chapter of the Association for Computational Linguistics:*  
640 *Volume 2, Short Papers*, pp. 8–14, Valencia, Spain, April  
641 2017. Association for Computational Linguistics. URL  
642 <https://aclanthology.org/E17-2002/>.
- 643 Liu, A., Pan, L., Hu, X., Meng, S., and Wen, L. A semantic  
644 invariant robust watermark for large language models. In  
645 *The Twelfth International Conference on Learning Rep-*  
646 *resentations*, 2024a. URL <https://openreview.net/forum?id=6p8lpe4MNF>.
- 647 Liu, A., Pan, L., Lu, Y., Li, J., Hu, X., Zhang, X., Wen,  
648 L., King, I., Xiong, H., and Yu, P. S. A survey of text  
649 watermarking in the era of large language models, 2024b.  
650 URL <https://arxiv.org/abs/2312.07913>.
- 651 Liu, Y. and Bu, Y. Adaptive text watermark for large lan-  
652 guage models, 2024. URL <https://arxiv.org/abs/2401.13927>.
- 653 Lu, Y., Liu, A., Yu, D., Li, J., and King, I. An entropy-  
654 based text watermarking detection method, 2024a. URL  
655 <https://arxiv.org/abs/2403.13485>.
- 656 Lu, Y., Zhu, W., Li, L., Qiao, Y., and Yuan, F. LLa-  
657 MAX: Scaling linguistic horizons of LLM by enhanc-  
658 ing translation capabilities beyond 100 languages. In  
659 Al-Onaizan, Y., Bansal, M., and Chen, Y.-N. (eds.),  
*Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 10748–10772, Miami, Florida, USA, November 2024b. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.631. URL <https://aclanthology.org/2024.findings-emnlp.631>.
- Luo, Y., Lin, K., Gu, C., Hou, J., Wen, L., and Ping, L. Lost in overlap: Exploring logit-based watermark collision in LLMs. In Chiruzzo, L., Ritter, A., and Wang, L. (eds.), *Findings of the Association for Computational Linguistics: NAACL 2025*, pp. 620–637, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN 979-8-89176-195-7. doi: 10.18653/v1/2025.findings-naacl.37. URL <https://aclanthology.org/2025.findings-naacl.37/>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library, 2019. URL <https://arxiv.org/abs/1912.01703>.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer, 2023. URL <https://arxiv.org/abs/1910.10683>.
- Ren, J., Xu, H., Liu, Y., Cui, Y., Wang, S., Yin, D., and Tang, J. A robust semantics-based watermark for large language model against paraphrasing, 2024. URL <https://arxiv.org/abs/2311.08721>.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. Transformers: State-of-the-art natural language processing. In Liu, Q. and Schlangen, D. (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October

660 2020. Association for Computational Linguistics. doi:  
661 10.18653/v1/2020.emnlp-demos.6. URL [https://](https://aclanthology.org/2020.emnlp-demos.6/)  
662 [aclanthology.org/2020.emnlp-demos.6/](https://aclanthology.org/2020.emnlp-demos.6/).

663  
664 Wu, Y., Hu, Z., Guo, J., Zhang, H., and Huang, H. A  
665 resilient and accessible distribution-preserving water-  
666 mark for large language models, 2024. URL [https:](https://arxiv.org/abs/2310.07710)  
667 [//arxiv.org/abs/2310.07710](https://arxiv.org/abs/2310.07710).

668 Zhao, X., Ananth, P., Li, L., and Wang, Y.-X. Provable  
669 robust watermarking for ai-generated text, 2023. URL  
670 <https://arxiv.org/abs/2306.17439>.  
671

672  
673  
674  
675  
676  
677  
678  
679  
680  
681  
682  
683  
684  
685  
686  
687  
688  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701  
702  
703  
704  
705  
706  
707  
708  
709  
710  
711  
712  
713  
714

## Appendix

The appendices contain the following sections:

- Appendix A details the experimental settings,
- Appendix B contains additional experimental results,
- Appendix C contains our usage of AI assistants.
- Appendix D contains the details of our artifacts.

For transparency and reproducibility, our code is available at this anonymized URL<sup>1</sup>: <https://anonymous.4open.science/r/STEAM-BBFA>

### A. Experimental Setting

#### A.1. Hyperparameters

To ensure reproducibility, we detail the hyperparameters used for both the neural network training and the watermark generation/detection phases of our experiments.

**X-SIR neural network training.** The neural network component of X-SIR, which inherits its architecture from SIR, was trained using the following hyperparameters:

- Architecture. The model consists of 4 layers, with an input dimension of 1024, a hidden dimension of 500, and an output dimension of 300.
- Optimization. We used *Stochastic Gradient Descent (SGD)* with a *learning rate* of 0.006 and a *weight decay* of 0.2. A `StepLR` scheduler with a *step size* of 200 and a *gamma* of 0.1 was employed to adjust the learning rate during training.
- Training. The model was trained for 2000 epochs with a batch size of 32.

**Watermarking scheme parameters.** For the watermark generation and detection phases, the following parameters were used for each scheme:

- KGW. We used the default parameters recommended by Kirchenbauer et al. (2023): a green list proportion (*gamma*) of 0.25, a logit bias (*delta*) of 2.0, and the `minhash` seeding scheme.
- X-KGW. To create a direct comparison with XSIR, we set the context width to 1. The *gamma* and *delta* values were kept consistent with KGW at 0.25 and 2.0, respectively.
- X-SIR. We followed the original implementation, setting the window size to 5, the chunk size to 10, and the logit bias (*delta*) to 1.0. The multilingual sentence embeddings were generated using the `paraphrase-multilingual-mpnet-base-v2` model.

#### A.2. Computational Resources & Softwares

All experiments were conducted on a Google Cloud Platform instance of type `n1-standard-4`, equipped with 4 vCPUs, 15 GB of RAM, and two NVIDIA T4 GPUs.

All translations use the Google Translate service accessed through the `deep_translator` Python library, which provides a unified interface to various translation APIs. The translator mismatch experiment in §5.3 employs the DeepSeek API for back-translation.

We used Pytorch as our deep learning framework (Paszke et al., 2019), with CUDA support for GPU acceleration. In addition, we employed Hugging Face Transformers library<sup>2</sup> (Wolf et al., 2020) to access pretrained models and tokenizers.

<sup>1</sup>Anonymised repository and URL for double-blind review

<sup>2</sup><https://huggingface.co/>

### A.3. Description of X-KGW

X-KGW (Cross-lingual KGW) is a hybrid watermarking approach we introduce to combine the hash-based mechanism of KGW with the semantic clustering strategy of X-SIR. Unlike KGW, which partitions individual tokens, X-KGW operates at the cluster level. The process consists of three distinct phases:

1. **Semantic cluster construction.** Following the X-SIR framework, we first construct a multilingual semantic graph using bilingual translation dictionaries. The Louvain community detection algorithm is then applied to partition the vocabulary  $\mathcal{V}$  into  $\mathcal{C}$  disjoint semantic clusters, yielding a token-to-cluster mapping  $m : \mathcal{V} \rightarrow 0, 1, \dots, \mathcal{C} - 1$ .
2. **Hash-based cluster partitioning.** During text generation, at each timestep  $t$ , a context window of preceding tokens  $(w_{t-h}, \dots, w_{t-1})$  is used to compute a hash-based seed. This seed is then employed to pseudo-randomly partition the  $\mathcal{C}$  clusters into green and red sets, with a fraction  $\gamma$  designated as green.
3. **Cluster-based logit modification.** Finally, a positive bias  $\delta$  is applied to the logits of all tokens belonging to clusters assigned to the green set. The model then samples the next token from this modified probability distribution.

By combining KGW logit biasing with semantic clustering, X-KGW seeks to preserve watermark robustness under multilingual transformations while maintaining detection accuracy.

### A.4. Multilingual Dictionaries & Language Categorization

To construct our multilingual dictionary, we relied on the MUSE dictionary (Conneau et al., 2017), the same resource used by He et al. (2024) to build the semantic clusters. In addition to its role in dictionary-based clustering, we used MUSE to categorize the 17 languages included in our evaluation of §4.2, §4.3, §5.2, §5.3.

- A language was marked *high-resource* if it possesses extensive, non-English-centric dictionary mappings (i.e., bidirectional dictionaries with multiple other languages in the set).
- In contrast, languages whose resources are primarily English-centric, where MUSE provides only bidirectional dictionaries with English, were classified as either *medium-resource* or *low-resource*. The distinction between these two groups was determined by the size (i.e., the number of word pairs) of their respective English dictionaries.

### A.5. Definitions of Multilingual Watermarking Comparison Criteria

For clarity, we provide the definitions of the criteria used in Table 1 to compare multilingual watermarking techniques:

- **Multilingual support.** Designed to resist translation attacks.
- **Non-invasive.** Supporting multilingual does not change the logits during generation, so the text quality is guaranteed to be preserved.
- **Watermark-agnostic.** Can be combined with any watermarking technique without modification.
- **Tokenizer-agnostic.** Robustness against translation attacks does not depend on the tokenizer.
- **Medium/low-resource.** Robust against translation attacks to medium-/low-resource languages.
- **Retroactive support.** Allows adding new languages without regenerating the watermark key (red/green tokens split). Already generated texts can be detected in the new languages.

### A.6. Prompt for DeepSeek-V3.2-Exp Translation

To use DeepSeek-V3.2-Exp as a translation engine, we designed a structured prompt format. We define:

```
Source language: {src_lang}
Target language: {tgt_lang}
Input text: {response}
```

`src_lang` and `tgt_lang` indicate the source and target language codes. We convert these codes into their full language names using the `Language` class from the `langcodes`<sup>3</sup> library:

```
Language.make(language=src_lang).display_name()
Language.make(language=tgt_lang).display_name()
```

The final prompt provided to DeepSeek-V3.2-Exp:

```
Translate the following {Language.make(language=src_lang).display_name()}
text to {Language.make(language=tgt_lang).display_name()}:

{response}
```

### A.7. Back-Translation Candidate Languages

STEAM searches over a pool of candidate back-translation languages using Google Translate as the default translation service. Table 7 lists all candidate languages used in our experiments for § 5.2 § 5.3. These languages span diverse language families, scripts, and resource levels, covering high-resource (e.g., English, French, German), medium-resource (e.g., Hungarian, Romanian, Thai), and low-resource (e.g., Tigrinya, Xhosa, Yoruba) settings.

### A.8. Bayesian Optimisation Details

**Surrogate model.** We use a single-task Gaussian Process (GP) with a constant mean function as implemented by `SingleTaskGP` in BoTorch (Balandat et al., 2020). The GP hyperparameters (kernel lengthscales, output scale, and noise variance) are optimised by maximising the exact marginal log-likelihood using GPyTorch (Gardner et al., 2021). The GP is refitted from scratch at each BO iteration.

**Acquisition function.** We use Log Expected Improvement (`LogExpectedImprovement` in BoTorch), a numerically stable variant of expected improvement that operates in log-space to avoid vanishing gradients when the current best value is far from the predictive mean. Since the search space is a finite set of discrete languages rather than a continuous domain, we evaluate the acquisition function at all unevaluated candidate feature vectors and select the one with the highest value.

**Language feature vectors.** Each candidate probe language is represented by a 131-dimensional feature vector obtained by concatenating `syntax_knn` (103 dimensions) and `phonology_knn` (28 dimensions) from the URIEL typological database (Khan et al., 2025), accessed via `lang2vec` (Littell et al., 2017). Feature vectors are pre-computed once and reused across all texts.

**Budget.** Each text is allocated a maximum of 20 evaluations (3 initial + 17 BO iterations).

**Implementation.** The full pipeline is implemented in Python using BoTorch 0.11+, GPyTorch 1.11+, and PyTorch 2.0+.

<sup>3</sup><https://pypi.org/project/langcodes/>

Table 7. All 133 candidate back-translation languages supported by STEAM.

Code	Language	Code	Language	Code	Language
af	Afrikaans	hmn	Hmong	qu	Quechua
sq	Albanian	hu	Hungarian	ro	Romanian
am	Amharic	is	Icelandic	ru	Russian
ar	Arabic	ig	Igbo	sm	Samoan
hy	Armenian	ilo	Ilocano	sa	Sanskrit
as	Assamese	id	Indonesian	gd	Scottish Gaelic
ay	Aymara	ga	Irish	nso	Sepedi
az	Azerbaijani	it	Italian	sr	Serbian
bm	Bambara	ja	Japanese	st	Sesotho
eu	Basque	jw	Javanese	sn	Shona
be	Belarusian	kn	Kannada	sd	Sindhi
bn	Bengali	kk	Kazakh	si	Sinhala
bho	Bhojpuri	km	Khmer	sk	Slovak
bs	Bosnian	rw	Kinyarwanda	sl	Slovenian
bg	Bulgarian	gom	Konkani	so	Somali
ca	Catalan	ko	Korean	es	Spanish
ceb	Cebuano	kri	Krio	su	Sundanese
ny	Chichewa	ku	Kurdish (Kurmanji)	sw	Swahili
zh-CN	Chinese (Simp.)	ckb	Kurdish (Sorani)	sv	Swedish
zh-TW	Chinese (Trad.)	ky	Kyrgyz	tg	Tajik
co	Corsican	lo	Lao	ta	Tamil
hr	Croatian	la	Latin	tt	Tatar
cs	Czech	lv	Latvian	te	Telugu
da	Danish	ln	Lingala	th	Thai
dv	Dhivehi	lt	Lithuanian	ti	Tigrinya
doi	Dogri	lg	Luganda	ts	Tsonga
nl	Dutch	lb	Luxembourgish	tr	Turkish
en	English	mk	Macedonian	tk	Turkmen
eo	Esperanto	mai	Maithili	ak	Twi
et	Estonian	mg	Malagasy	uk	Ukrainian
ee	Ewe	ms	Malay	ur	Urdu
tl	Filipino	ml	Malayalam	ug	Uyghur
fi	Finnish	mt	Maltese	uz	Uzbek
fr	French	mi	Maori	vi	Vietnamese
fy	Frisian	mr	Marathi	cy	Welsh
gl	Galician	mni-Mtei	Meitei	xh	Xhosa
ka	Georgian	lus	Mizo	yi	Yiddish
de	German	mn	Mongolian	yo	Yoruba
el	Greek	my	Myanmar	zu	Zulu
gn	Guarani	ne	Nepali	gu	Gujarati
no	Norwegian	ht	Haitian Creole	or	Odia
ha	Hausa	om	Oromo	haw	Hawaiian
ps	Pashto	iw	Hebrew	fa	Persian
hi	Hindi	pl	Polish	pt	Portuguese
pa	Punjabi				

## B. Additional Results

### B.1. Hold-out languages for X-SIR

Table 8. **Semantic clustering (XSIR) is weak for hold-out unsupported languages.**  $\Delta$  measures the robustness gains against a translation attack on a language after it has been supported by XSIR. The semantic clustering of tokens is applied on all the original five languages of XSIR (en, fr, de, zh, ja) for *supported*, and on all but the held-out language for *held-out*. Aya-23 8B generates a text in the *Prompt* language, then the translation attack is applied on the held-out language. Red indicates that XSIR performs worst after supporting the held-out language.

Languages		AUC ( $\uparrow$ )			TPR@1% ( $\uparrow$ )		
Held-out	Prompt	Held-Out	Supported	$\Delta$	Held-Out	Supported	$\Delta$
en	fr	0.795 $\pm$ 0.045	0.816 $\pm$ 0.014	+0.021	0.198 $\pm$ 0.049	0.149 $\pm$ 0.042	-0.049
	de	0.780 $\pm$ 0.054	0.811 $\pm$ 0.018	+0.031	0.172 $\pm$ 0.047	0.168 $\pm$ 0.039	-0.004
	zh	0.731 $\pm$ 0.020	0.669 $\pm$ 0.042	-0.062	0.141 $\pm$ 0.011	0.083 $\pm$ 0.014	-0.058
fr	en	0.757 $\pm$ 0.022	0.799 $\pm$ 0.027	+0.042	0.157 $\pm$ 0.016	0.139 $\pm$ 0.037	-0.018
	de	0.723 $\pm$ 0.020	0.781 $\pm$ 0.025	+0.058	0.101 $\pm$ 0.029	0.156 $\pm$ 0.061	+0.055
	zh	0.651 $\pm$ 0.026	0.638 $\pm$ 0.037	-0.013	0.076 $\pm$ 0.021	0.052 $\pm$ 0.020	-0.024
de	en	0.736 $\pm$ 0.011	0.802 $\pm$ 0.020	+0.067	0.153 $\pm$ 0.050	0.214 $\pm$ 0.035	+0.061
	fr	0.765 $\pm$ 0.004	0.784 $\pm$ 0.020	+0.019	0.139 $\pm$ 0.068	0.118 $\pm$ 0.039	-0.021
	zh	0.667 $\pm$ 0.041	0.642 $\pm$ 0.014	-0.025	0.073 $\pm$ 0.032	0.065 $\pm$ 0.016	-0.008
zh	en	0.644 $\pm$ 0.052	0.692 $\pm$ 0.041	+0.048	0.120 $\pm$ 0.046	0.111 $\pm$ 0.011	-0.009
	fr	0.671 $\pm$ 0.072	0.714 $\pm$ 0.045	+0.043	0.112 $\pm$ 0.053	0.069 $\pm$ 0.025	-0.043
	de	0.675 $\pm$ 0.048	0.701 $\pm$ 0.026	+0.026	0.105 $\pm$ 0.053	0.107 $\pm$ 0.027	+0.002
ja	en	0.685 $\pm$ 0.059	0.656 $\pm$ 0.017	-0.029	0.113 $\pm$ 0.024	0.070 $\pm$ 0.008	-0.043
	fr	0.698 $\pm$ 0.038	0.670 $\pm$ 0.018	-0.028	0.101 $\pm$ 0.033	0.089 $\pm$ 0.023	-0.012
	de	0.688 $\pm$ 0.037	0.669 $\pm$ 0.027	-0.019	0.138 $\pm$ 0.031	0.079 $\pm$ 0.005	-0.059
	zh	0.681 $\pm$ 0.043	0.658 $\pm$ 0.003	-0.023	0.110 $\pm$ 0.046	0.093 $\pm$ 0.016	-0.017

Languages		AUC ( $\uparrow$ )			TPR@1% ( $\uparrow$ )		
Held-out	Prompt	Held-Out	Supported	$\Delta$	Held-Out	Supported	$\Delta$
en	fr	0.901 $\pm$ 0.017	0.907 $\pm$ 0.015	+0.006	0.363 $\pm$ 0.047	0.275 $\pm$ 0.007	-0.088
	de	0.884 $\pm$ 0.043	0.894 $\pm$ 0.041	+0.010	0.379 $\pm$ 0.131	0.331 $\pm$ 0.043	-0.048
	zh	0.795 $\pm$ 0.043	0.827 $\pm$ 0.013	+0.032	0.450 $\pm$ 0.040	0.407 $\pm$ 0.011	-0.043
fr	en	0.743 $\pm$ 0.050	0.682 $\pm$ 0.026	-0.061	0.113 $\pm$ 0.056	0.068 $\pm$ 0.012	-0.045
	de	0.787 $\pm$ 0.050	0.687 $\pm$ 0.011	-0.100	0.161 $\pm$ 0.038	0.093 $\pm$ 0.007	-0.068
	zh	0.678 $\pm$ 0.006	0.681 $\pm$ 0.007	+0.003	0.086 $\pm$ 0.027	0.083 $\pm$ 0.005	-0.003
de	en	0.693 $\pm$ 0.017	0.692 $\pm$ 0.043	-0.001	0.069 $\pm$ 0.022	0.068 $\pm$ 0.025	-0.001
	fr	0.724 $\pm$ 0.023	0.725 $\pm$ 0.044	+0.001	0.098 $\pm$ 0.011	0.071 $\pm$ 0.009	-0.027
	zh	0.665 $\pm$ 0.018	0.693 $\pm$ 0.046	+0.028	0.065 $\pm$ 0.009	0.073 $\pm$ 0.014	+0.008
zh	en	0.666 $\pm$ 0.035	0.605 $\pm$ 0.012	-0.061	0.082 $\pm$ 0.052	0.026 $\pm$ 0.006	-0.056
	fr	0.704 $\pm$ 0.011	0.609 $\pm$ 0.024	-0.095	0.085 $\pm$ 0.041	0.036 $\pm$ 0.012	-0.049
	de	0.703 $\pm$ 0.022	0.636 $\pm$ 0.017	-0.067	0.088 $\pm$ 0.010	0.050 $\pm$ 0.008	-0.038
ja	en	0.576 $\pm$ 0.052	0.573 $\pm$ 0.024	-0.003	0.039 $\pm$ 0.019	0.033 $\pm$ 0.005	-0.006
	fr	0.630 $\pm$ 0.066	0.581 $\pm$ 0.041	-0.049	0.067 $\pm$ 0.016	0.025 $\pm$ 0.014	-0.042
	de	0.624 $\pm$ 0.055	0.589 $\pm$ 0.039	-0.035	0.074 $\pm$ 0.023	0.037 $\pm$ 0.018	-0.037
	zh	0.663 $\pm$ 0.059	0.650 $\pm$ 0.026	-0.013	0.121 $\pm$ 0.065	0.075 $\pm$ 0.040	-0.046

Table 9. **Semantic clustering (XSIR) performs poorly on hold-out unsupported languages.**  $\Delta$  measures the robustness gains against a translation attack on a language after it has been supported by XSIR. The semantic clustering of tokens is applied on all the original five languages of XSIR (en, fr, de, zh, ja) for *supported*, and on all but the held-out language for *held-out*. LLaMA-3.2 1B generates a text in the *Prompt* language, then the translation attack is applied on the held-out language. Red indicates that XSIR performs worst after supporting the held-out language.

## B.2. Unsupported languages for X-SIR & X-KGW

Table 10. **Semantic clustering is weak for unsupported languages.** Watermark strength (AUC and TPR@1%) of X-SIR and X-KGW, limited to the five originally supported languages (en, fr, de, zh, ja). Aya-23 8B generates English text, which is then translated into a new unsupported language for evaluation. Minimum marks the weakest robustness (best attack case).

New Lang.	X-SIR ( $\uparrow$ )		X-KGW ( $\uparrow$ )	
	AUC	TPR@1%	AUC	TPR@1%
it	0.796	0.177	0.772	0.238
es	0.754	0.155	0.807	0.230
pt	0.775	0.133	0.792	0.286
pl	0.749	0.127	0.762	0.236
nl	0.776	0.164	0.808	0.314
hr	0.726	0.124	0.757	0.210
cs	0.773	0.111	0.754	0.254
da	0.734	0.161	0.764	0.266
ko	0.729	0.136	0.754	0.226
ar	0.687	0.093	0.765	0.168
Min.	0.687 (ar)	0.093 (ar)	0.754 (cs, ko)	0.168 (ar)

New Lang.	X-SIR ( $\uparrow$ )		X-KGW ( $\uparrow$ )	
	AUC	TPR@1%	AUC	TPR@1%
it	0.699	0.069	0.760	0.212
es	0.665	0.076	0.744	0.222
pt	0.641	0.059	0.722	0.152
pl	0.679	0.069	0.677	0.144
nl	0.754	0.095	0.781	0.244
hr	0.660	0.066	0.733	0.162
cs	0.650	0.064	0.759	0.190
da	0.675	0.093	0.765	0.196
ko	0.673	0.062	0.672	0.124
ar	0.655	0.055	0.704	0.168
Min.	0.641 (pt)	0.055 (ar)	0.672 (ko)	0.124 (ko)

Table 11. **Semantic clustering is weak for unsupported languages.** Watermark strength (AUC and TPR@1%) of X-SIR and X-KGW, limited to the five originally supported languages (en, fr, de, zh, ja). LLaMA-3.2 1B generates English text, which is then translated into a new unsupported language for evaluation. Minimum marks the weakest robustness (best attack case).

New Lang.	X-SIR ( $\uparrow$ )		X-KGW ( $\uparrow$ )	
	AUC	TPR@1%	AUC	TPR@1%
it	0.829	0.335	0.860	0.510
es	0.810	0.314	0.864	0.490
pt	0.812	0.337	0.861	0.498
pl	0.812	0.308	0.817	0.420
nl	0.845	0.351	0.896	0.584
hr	0.804	0.279	0.822	0.344
cs	0.798	0.305	0.838	0.444
da	0.832	0.357	0.871	0.436
ko	0.792	0.276	0.800	0.390
ar	0.765	0.251	0.815	0.358
Min.	0.765 (ar)	0.251 (ar)	0.800 (ko)	0.344 (hr)

Table 12. **Semantic clustering is weak for unsupported languages.** Watermark strength (AUC and TPR@1%) of X-SIR and X-KGW, limited to the five originally supported languages (en, fr, de, zh, ja). LLaMAX-3 8B generates English text, which is then translated into a new unsupported language for evaluation. Minimum marks the weakest robustness (best attack case).

CWRA Attack New Language	Aya-23 8B ( $\uparrow$ )		LLaMA-3.2 1B ( $\uparrow$ )		LLaMAX-3 8B ( $\uparrow$ )	
	AUC	TPR@1%	AUC	TPR@1%	AUC	TPR@1%
it	0.746	0.194	0.855	0.245	0.826	0.313
es	0.751	0.147	0.830	0.217	0.816	0.319
pt	0.781	0.179	0.854	0.244	0.827	0.336
pl	0.793	0.195	0.859	0.289	0.815	0.299
nl	0.836	0.252	0.900	0.375	0.835	0.360
hr	0.810	0.236	0.853	0.269	0.790	0.291
cs	0.785	0.180	0.835	0.201	0.787	0.309
da	0.857	0.247	0.864	0.243	0.830	0.331
ko	0.750	0.157	0.852	0.255	0.809	0.309
ar	0.704	0.209	0.822	0.222	0.771	0.286
Minimum	0.704 (ar)	0.147 (es)	0.822 (ar)	0.201 (cs)	0.771 (ar)	0.286 (ar)

Table 13. **Semantic clustering (XSIR) performs inconsistently on an expanded set of supported languages.** The semantic clustering is applied using an expanded set of 17 newly supported languages. A prompt in English is first translated into each target language. Aya-23 8B, LLaMA-3.2 1B, and LLaMAX-3 8B are then prompted with the translated input to generate text in the target language. Finally, the CWRA attack is applied by translating the generated text back into English. Baseline is the average on the original supported languages. Higher values indicate better robustness. Minimum indicates the worst-case robustness, i.e., the best language for an attack.

## B.3. Supported languages for X-SIR &amp; X-KGW

Translation Attack		X-SIR ( $\uparrow$ )		X-KGW ( $\uparrow$ )	
Type	Language	AUC	TPR@1%	AUC	TPR@1%
High-resource	fr	0.702	0.085	0.719	0.166
	de	0.708	0.067	0.752	0.186
	it	0.712	0.111	0.750	0.230
	es	0.703	0.089	0.724	0.222
	pt	0.726	0.102	0.747	0.206
Medium-resource	pl	0.657	0.065	0.703	0.188
	nl	0.722	0.091	0.787	0.252
	ru	0.635	0.075	0.656	0.100
	hi	0.611	0.037	0.620	0.084
	ko	0.673	0.055	0.701	0.154
Low-resource	ja	0.571	0.042	0.598	0.110
	bn	0.825	0.509	0.701	0.078
	fa	0.584	0.055	0.673	0.086
	vi	0.691	0.084	0.722	0.186
	iw	0.622	0.026	0.702	0.134
	uk	0.613	0.064	0.725	0.118
	ta	0.749	0.095	0.672	0.108
Minimum		0.571 (ja)	0.026 (iw)	0.598 (ja)	0.078 (bn)

Table 14. **Semantic clustering performs poorly on an expanded set of supported languages.** The semantic clustering is applied using an expanded set of 17 newly supported languages. LLaMA-3.2 1B generates a text in English, then the translation attack is applied using each of these supported languages as target language. Higher values indicate better robustness. Minimum indicates the worst-case robustness, i.e., the best language for an attack.

Translation Attack		X-SIR ( $\uparrow$ )		X-KGW ( $\uparrow$ )	
Type	Language	AUC	TPR@1%	AUC	TPR@1%
High-resource	fr	0.804	0.249	0.852	0.466
	de	0.833	0.399	0.850	0.484
	it	0.829	0.336	0.870	0.478
	es	0.811	0.319	0.869	0.506
	pt	0.726	0.338	0.863	0.454
Medium-resource	pl	0.812	0.308	0.847	0.410
	nl	0.847	0.355	0.882	0.592
	ru	0.787	0.256	0.821	0.368
	hi	0.702	0.215	0.714	0.228
	ko	0.792	0.276	0.822	0.422
Low-resource	ja	0.714	0.187	0.705	0.206
	bn	0.588	0.086	0.765	0.244
	fa	0.755	0.268	0.829	0.398
	vi	0.772	0.238	0.802	0.328
	iw	0.719	0.196	0.808	0.444
uk	0.794	0.309	0.817	0.118	
ta	0.561	0.067	0.789	0.316	
Minimum		0.561 (ta)	0.067 (ta)	0.705 (ja)	0.118 (uk)

Table 15. **Semantic clustering performs poorly on an expanded set of supported languages.** The semantic clustering is applied using an expanded set of 17 newly supported languages. LLaMAX-3 8B generates a text in English, then the translation attack is applied using each of these supported languages as target language. Higher values indicate better robustness. Minimum indicates the worst-case robustness, i.e., the best language for an attack.

CWRA Attack		Aya-23 8B ( $\uparrow$ )		LLaMA-3.2 1B ( $\uparrow$ )		LLaMAX-3 8B ( $\uparrow$ )	
Type	Language	AUC	TPR@1%	AUC	TPR@1%	AUC	TPR@1%
High-resource	fr	0.831	0.201	0.898	0.421	0.845	0.345
	de	0.820	0.198	0.902	0.394	0.841	0.361
	it	0.817	0.196	0.872	0.331	0.825	0.315
	es	0.819	0.200	0.859	0.281	0.816	0.320
	pt	0.801	0.191	0.876	0.348	0.827	0.335
Medium-resource	pl	0.804	0.202	0.881	0.381	0.815	0.291
	nl	0.859	0.265	0.899	0.434	0.834	0.367
	ru	0.769	0.175	0.824	0.226	0.771	0.233
	hi	0.710	0.147	0.771	0.304	0.744	0.263
	ko	0.769	0.178	0.854	0.340	0.805	0.299
Low-resource	ja	0.787	0.278	0.904	0.644	0.720	0.314
	bn	0.721	0.255	0.934	0.628	0.784	0.354
	fa	0.691	0.113	0.796	0.260	0.733	0.246
	vi	0.781	0.180	0.865	0.325	0.805	0.301
	iw	0.726	0.124	0.815	0.207	0.729	0.245
uk	0.769	0.157	0.849	0.253	0.750	0.225	
ta	0.819	0.405	0.917	0.516	0.740	0.391	
Minimum		0.691 (fa)	0.113 (fa)	0.771 (hi)	0.207 (iw)	0.720 (ja)	0.225 (uk)

Table 16. **Semantic clustering (XSIR) performs inconsistently on an expanded set of supported languages.** The semantic clustering is applied using an expanded set of 17 newly supported languages. A prompt in English is first translated into each target language. Aya-23 8B, LLaMA-3.2 1B, and LLaMAX-3 8B are then prompted with the translated input to generate text in the target language. Finally, the CWRA attack is applied by translating the generated text back into English. Higher values indicate better robustness. Minimum indicates the worst-case robustness, i.e., the best language for an attack.

## B.4. STEAM

Translation Attack		AUC ( $\uparrow$ )				TPR@1% ( $\uparrow$ )			
Type	Language	KGW	X-KGW	X-SIR	STEAM $\hat{C}_m$	KGW	X-KGW	X-SIR	STEAM $\hat{C}_m$
High-resource	fr	0.655	0.719	0.702	<b>0.969</b>	0.052	0.166	0.085	<b>0.618</b>
	de	0.649	0.752	0.708	<b>0.975</b>	0.090	0.186	0.067	<b>0.530</b>
	it	0.617	0.750	0.712	<b>0.977</b>	0.118	0.230	0.111	<b>0.620</b>
	es	0.616	0.724	0.703	<b>0.969</b>	0.122	0.222	0.089	<b>0.522</b>
	pt	0.651	0.747	0.726	<b>0.970</b>	0.096	0.206	0.102	<b>0.476</b>
Medium-resource	pl	0.622	0.703	0.657	<b>0.975</b>	0.088	0.188	0.065	<b>0.526</b>
	nl	0.719	0.787	0.722	<b>0.973</b>	0.126	0.252	0.091	<b>0.580</b>
	ru	0.629	0.656	0.635	<b>0.970</b>	0.052	0.100	0.075	<b>0.482</b>
	hi	0.568	0.620	0.611	<b>0.974</b>	0.048	0.084	0.037	<b>0.598</b>
	ko	0.625	0.701	0.673	<b>0.971</b>	0.068	0.154	0.055	<b>0.416</b>
ja	0.578	0.598	0.571	<b>0.973</b>	0.048	0.110	0.042	<b>0.482</b>	
Low-resource	bn	0.574	0.701	0.825	<b>0.967</b>	0.020	0.078	<b>0.509</b>	0.496
	fa	0.586	0.673	0.584	<b>0.979</b>	0.082	0.086	0.055	<b>0.664</b>
	vi	0.658	0.722	0.691	<b>0.975</b>	0.082	0.186	0.084	<b>0.458</b>
	iw	0.495	0.702	0.622	<b>0.966</b>	0.042	0.134	0.026	<b>0.502</b>
	uk	0.629	0.725	0.613	<b>0.973</b>	0.084	0.118	0.064	<b>0.482</b>
ta	0.877	0.672	0.749	<b>0.977</b>	0.272	0.108	0.095	<b>0.460</b>	

Table 17. STEAM  $\hat{C}_m$  is consistently better than semantic clustering by a large margin. Watermark strength (AUC and TPR@1%) of multilingual watermarking techniques with 17 supported languages and LLaMA-3.2 1B. Red indicates that the defence reduces robustness (lower than the undefended KGW baseline). Bolded is best.

New Lang.	AUC ( $\uparrow$ )				TPR@1% ( $\uparrow$ )			
	KGW	X-KGW	X-SIR	STEAM $\hat{C}_m$	KGW	X-KGW	X-SIR	STEAM $\hat{C}_m$
it	0.733	0.772	0.796	<b>0.966</b>	0.202	0.238	0.177	<b>0.494</b>
es	0.717	0.807	0.754	<b>0.967</b>	0.232	0.230	0.155	<b>0.548</b>
pt	0.732	0.792	0.775	<b>0.971</b>	0.242	0.286	0.133	<b>0.466</b>
pl	0.730	0.762	0.749	<b>0.960</b>	0.248	0.236	0.127	<b>0.344</b>
nl	0.768	0.808	0.776	<b>0.966</b>	0.286	0.314	0.164	<b>0.358</b>
hr	0.706	0.757	0.726	<b>0.965</b>	0.194	0.210	0.124	<b>0.362</b>
cs	0.717	0.754	0.773	<b>0.974</b>	0.212	0.254	0.111	<b>0.554</b>
da	0.713	0.764	0.734	<b>0.971</b>	0.196	0.266	0.161	<b>0.448</b>
ko	0.732	0.754	0.729	<b>0.961</b>	0.220	0.226	0.136	<b>0.322</b>
ar	0.689	0.765	0.687	<b>0.971</b>	0.186	0.168	0.093	<b>0.588</b>

Table 18. STEAM  $\hat{C}_m$  performs on par with other multilingual methods on unsupported languages. Watermark strength (AUC and TPR@1%) of multilingual watermarking techniques with 10 unsupported languages and Aya-23 8B. Red indicates that the defence reduces robustness (lower than the undefended KGW baseline). Bolded is best




New Lang.	AUC ( $\uparrow$ )				TPR@1% ( $\uparrow$ )			
	KGW	X-KGW	X-SIR	STEAM 	KGW	X-KGW	X-SIR	STEAM 
it	0.620	0.760	0.699	<b>0.964</b>	0.108	0.212	0.069	<b>0.350</b>
es	0.616	0.744	0.665	<b>0.961</b>	0.122	0.222	0.076	<b>0.358</b>
pt	0.652	0.722	0.641	<b>0.966</b>	0.096	0.152	0.059	<b>0.408</b>
pl	0.617	0.677	0.679	<b>0.965</b>	0.088	0.144	0.069	<b>0.310</b>
nl	0.714	0.781	0.754	<b>0.957</b>	0.112	0.244	0.095	<b>0.258</b>
hr	0.611	0.733	0.660	<b>0.966</b>	0.078	0.162	0.066	<b>0.448</b>
cs	0.655	0.759	0.650	<b>0.956</b>	0.072	0.190	0.064	<b>0.228</b>
da	0.655	0.765	0.675	<b>0.962</b>	0.080	0.196	0.093	<b>0.346</b>
ko	0.623	0.672	0.673	<b>0.961</b>	0.066	0.124	0.062	<b>0.300</b>
ar	0.635	<b>0.704</b>	0.655	0.670	0.110	<b>0.168</b>	0.055	0.078

Table 19. STEAM  outperforms semantic clustering methods for unsupported languages. Watermark strength (AUC and TPR@1%) of multilingual watermarking techniques with 10 unsupported languages and LLaMA-3.2 1B. Bold marks the best per row; red indicates a defended score lower than the KGW baseline.




Two-Step Translation Attack		STEAM 	
Language 1	Language 2	AUC $\uparrow$	TPR@1% $\uparrow$
High-resource	None	<b>0.976</b>	<b>0.570</b>
	de	0.899	0.343
	ko	0.833	0.227
Medium-resource	bn	0.864	0.252
	None	<b>0.975</b>	<b>0.553</b>
	de	0.910	0.349
Low-resource	ko	0.890	0.239
	bn	0.872	0.396
	None	<b>0.976</b>	<b>0.586</b>
Low-resource	de	0.875	0.164
	ko	0.938	0.252
	bn	0.871	0.325

Table 20. STEAM  remains robust under multi-step attacks. Aya-23 8B generates text in English that is translated to the 17 supported languages (*Language 1*). A second translation step is then applied using *Language 2* to compute the AUC and TPR@1%. None indicates the single-step translation baseline.

## B.5. Inference Cost

STEAM introduces no additional cost at generation time, as text is produced by the unmodified watermarking scheme (e.g., vanilla KGW). The only cost is at detection time, where STEAM requires translating the suspect text into back-translation languages (Table 21). Google Translate is free, while high-quality LLM-based translators such as DeepSeek-V3.2-Exp cost \$0.003 per text, making STEAM practical even at scale.


Model	Cost per text (\$)
Google Translate	0.000
DeepSeek-V3.2-Exp	0.060
GPT-5 mini	0.180
Gemini 2.5 Flash	0.220

Table 21. STEAM  requires minimal cost at detection time. Total cost in US dollars for 20 back-translations per text (the maximum number of evaluated languages by BO). Google Translate is free. STEAM adds no cost at generation time.


## B.6. Text Length Analysis

We analyse the impact of text length on watermark detection. For each language, we split texts into three equal-sized bins (short, medium, and long) based on percentiles of token length (Table 22 and Table 23). STEAM maintains strong detection across all length categories, with average AUC above 0.97 across all text lengths for Aya-23 and LLaMA-3.2 models. Even in the most challenging case (Hebrew, long texts, LLaMA-3.2 1B), STEAM achieves an AUC of 0.899.

Translation Attack		AUC ( $\uparrow$ )			TPR@1% ( $\uparrow$ )		
Type	Language	Short	Medium	Long	Short	Medium	Long
High-resource	fr	0.994	0.971	0.951	0.882	0.561	0.319
	de	0.988	0.979	0.948	0.827	0.619	0.451
	it	0.994	0.984	0.953	0.835	0.582	0.317
	es	0.992	0.970	0.958	0.856	0.444	0.378
	pt	0.993	0.975	0.956	0.810	0.464	0.323
Medium-resource	pl	0.987	0.964	0.975	0.754	0.408	0.433
	nl	0.993	0.985	0.963	0.835	0.641	0.325
	ru	0.987	0.973	0.949	0.753	0.500	0.317
	hi	0.996	0.976	0.960	0.893	0.562	0.344
	ko	0.988	0.982	0.923	0.687	0.612	0.220
Low-resource	ja	0.988	0.984	0.951	0.717	0.653	0.242
	bn	0.989	0.979	0.963	0.789	0.575	0.250
	fa	0.993	0.980	0.958	0.856	0.549	0.494
	vi	0.987	0.983	0.959	0.789	0.643	0.282
	iw	0.990	0.970	0.956	0.826	0.503	0.372
uk	0.992	0.976	0.970	0.713	0.465	0.368	
ta	0.992	0.952	0.948	0.840	0.389	0.195	

Table 22. STEAM  maintains strong AUC across all three text lengths. Watermark strength (AUC and TPR@1%) for short, medium, and long texts (bottom, middle, and top third by token count) using Aya-23 8B.

Translation Attack		AUC ( $\uparrow$ )			TPR@1% ( $\uparrow$ )		
Type	Language	Short	Medium	Long	Short	Medium	Long
High-resource	fr	0.994	0.960	0.951	0.922	0.243	0.390
	de	0.996	0.963	0.957	0.861	0.471	0.146
	it	0.996	0.980	0.947	0.899	0.583	0.018
	es	0.991	0.981	0.930	0.826	0.653	0.110
	pt	0.996	0.957	0.945	0.922	0.254	0.146
Medium-resource	pl	0.995	0.966	0.977	0.916	0.227	0.704
	nl	0.996	0.970	0.945	0.922	0.524	0.183
	ru	0.998	0.938	0.956	0.946	0.220	0.429
	hi	0.994	0.980	0.941	0.916	0.639	0.061
	ko	0.988	0.957	0.979	0.711	0.335	0.616
Low-resource	ja	0.992	0.965	0.954	0.820	0.439	0.272
	bn	0.993	0.960	0.943	0.795	0.365	0.323
	fa	0.993	0.978	0.944	0.861	0.600	0.250
	vi	0.994	0.972	0.955	0.783	0.285	0.191
	iw	0.996	0.977	0.899	0.899	0.361	0.130
uk	0.998	0.958	0.949	0.904	0.278	0.177	
ta	0.994	0.969	0.961	0.837	0.269	0.172	

Table 23. STEAM  maintains strong AUC across all three text lengths. Watermark strength (AUC and TPR@1%) for short, medium, and long texts (bottom, middle, and top third by token count) using LLaMA-3.2 1B.

**B.7. Sub-character Token Distributions**

As discussed in §5.1, the z-score normalization component is designed to calibrate STEAM’s detection mechanism against statistical noise introduced by tokenizer limitations. Figures 5 and 6 show the token distribution for two severely affected low-resource languages, Bengali and Tamil.




Figure 5. Tokenization of low-resource languages creates highly concentrated sub-character tokens. Percentage of top 10 tokens for Bengali (a) and for Tamil (b) using Aya-23 8B.




Figure 6. Tokenization of low-resource languages creates highly concentrated sub-character tokens. Percentage of top 10 tokens for Bengali (a) and for Tamil (b) using LLaMA-3.2 1B.

## C. Usage of AI Assistants

For coding-related tasks, we relied on Claude 4.5 Sonnet and GitHub Copilot. We use GPT-5 and Claude for light editing (re-wording, grammar, proof-checking) to help writing the paper. For translation tasks in the experimental setting of §5.3, we use DeepSeek-V3.2-Exp and GPT-4o-mini as the translation models.

## D. Artifacts

### D.1. Artifacts License

All datasets, models, and code used in this work comply with their original licenses.

- MUSE Dictionary<sup>4</sup> (Conneau et al., 2017): Released under the Creative Commons Attribution–NonCommercial 4.0 International (CC BY-NC 4.0) license. Use is restricted to non-commercial research and requires attribution to the original authors.
- Aya-23 8B<sup>5</sup>: Released under (CC BY-NC 4.0) license.
- LLaMA-3.2 1B<sup>6</sup>: Released under the LLaMA 3.2 Community License Agreement. This license allows research and educational use but restricts commercial deployment without explicit permission from Meta.
- LLaMAX3 8B<sup>7</sup>: Released under the MIT License, which permits reuse, modification, and redistribution for both commercial and non-commercial purposes, provided that attribution and the original license terms are preserved.
- DeepSeek-V3.2-Exp<sup>8</sup>: Released under the MIT License.
- mC4 Dataset<sup>9</sup> (Raffel et al., 2023): Licensed under the Open Data Commons Attribution License (ODC-BY). This allows redistribution, reuse, and adaptation of the dataset, provided that appropriate credit is given.
- `deep_translator`<sup>10</sup> python package: Released under the MIT License.
- `openai`<sup>11</sup> python package: Released under the Apache License 2.0. This license permits use, modification, and redistribution for both commercial and non-commercial purposes

### D.2. Artifact Use Consistent With Intended Use

All datasets and models were used in line with their intended research purposes and licences. We used the mC4 dataset (Raffel et al., 2023) and open multilingual models (Aya-23-8B, LLaMA-3.2-1B, LLaMAX-8B) strictly for evaluation within academic settings. No data or model outputs were used for deployment or commercial applications. Our method STEAM is released only for research use and is compatible with the original access conditions of all components. No personal data were processed.

<sup>4</sup><https://github.com/facebookresearch/MUSE?tab=License-1-ov-file>

<sup>5</sup><https://huggingface.co/Coherelabs/aya-23-8B>

<sup>6</sup><https://huggingface.co/meta-llama/llama-3.2-1B/blob/main/LICENSE.txt>

<sup>7</sup><https://huggingface.co/LLaMAX/LLaMAX3-8B>

<sup>8</sup><https://huggingface.co/deepseek-ai/DeepSeek-V3.2-Exp>

<sup>9</sup><https://huggingface.co/datasets/allenai/c4>

<sup>10</sup><https://deep-translator.readthedocs.io/en/latest/README.html>

<sup>11</sup><https://pypi.org/project/openai/>