



AGMark: Attention-Guided Dynamic Watermarking for Large Vision Language Models

Yue Li*
East China Normal University
Shanghai, China
yue_li@stu.ecnu.edu.cn

Xin Yi*
East China Normal University
Shanghai, China
xinyi@stu.ecnu.edu.cn

Dongsheng Shi
East China Normal University
Shanghai, China
51275901023@stu.ecnu.edu.cn

Yongyi Cui
East China Normal University
Shanghai, China
yycui@stu.ecnu.edu.cn

Gerard de Melo
Hasso Plattner Institute
Potsdam, Germany
University of Potsdam
Potsdam, Germany
gdm@demelo.org

Linlin Wang†
East China Normal University
Shanghai, China
City University of Hong Kong
Hong Kong, China
llwang@cs.ecnu.edu.cn

Abstract

Watermarking has emerged as a pivotal solution for content traceability and intellectual property protection in large vision language models (LVLMs). However, vision-agnostic watermarks may introduce visually irrelevant tokens and disrupt visual grounding by enforcing indiscriminate pseudo-random biases. Additionally, current vision-specific watermarks rely on a static, one-time estimation of vision-critical weights and ignore the weight distribution density when determining the proportion of protected tokens. This design fails to account for dynamic changes in visual dependence during generation and may introduce low-quality tokens in the long tail. To address these challenges, we propose Attention-Guided Dynamic Watermarking (AGMark), a novel framework that embeds detectable signals while largely preserving visual-semantic fidelity. At each decoding step, AGMark first dynamically identifies semantic-critical evidence based on attention weights for visual relevance, together with context-aware coherence cues, resulting in a more adaptive and well-calibrated evidence-weight distribution. It then determines the proportion of semantic-critical tokens by jointly considering uncertainty awareness (token entropy) and evidence calibration (weight density), thereby enabling more reliable adaptive vocabulary partitioning to avoid irrelevant tokens. Empirical results consistently confirm that AGMark outperforms conventional methods, substantially improving generation quality and yielding particularly strong gains in visual semantic fidelity in the later stages of generation. Our framework maintains highly competitive detection performance (at least 99.36% AUC) and robust attack resilience (at least 88.61% AUC) without sacrificing inference efficiency, taking a significant step toward reliability-preserving multimodal watermarking.

*Both authors contributed equally to this research.

†Corresponding author.



This work is licensed under a Creative Commons Attribution 4.0 International License. *KDD 2026, Jeju Island, Republic of Korea.*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2259-2/2026/08

<https://doi.org/10.1145/3770855.3817700>

CCS Concepts

• Security and privacy → Digital rights management; • Computing methodologies → Computer vision; Natural language processing.

Keywords

Dynamic Watermarking; Large Vision Language Models

ACM Reference Format:

Yue Li, Xin Yi, Dongsheng Shi, Yongyi Cui, Gerard de Melo, and Linlin Wang. 2026. AGMark: Attention-Guided Dynamic Watermarking for Large Vision Language Models. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD 2026), August 9–13, 2026, Jeju Island, Republic of Korea*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3770855.3817700>

Resource Availability:

The source code of this paper has been made publicly available at <https://doi.org/10.5281/zenodo.20339442>.

1 Introduction

Large vision language models (LVLMs) demonstrate strong capabilities across a broad spectrum of downstream tasks [20, 34, 37]. However, their increasing deployment also introduces risks of malicious misuse, motivating the urgent need for robust intellectual property protection and reliable mechanisms to ensure content authenticity [8, 19]. Model watermarking addresses this by embedding identifiable signals into generated content for source attribution and traceability [8, 36]. Existing watermarking techniques for such models can be categorized into logits-based and sampling-based methods [25], the former becoming the dominant line of work [22]. The pioneering method KGW [14] employs a secret hash key to partition the vocabulary into green and red token sets and biases the generation process toward green tokens. Building on this framework, follow-up studies propose detection statistics based on token entropy [10, 16] or the green list probability mass [32], aiming to better balance watermark detectability and generation quality.

The majority of relevant watermarking techniques were developed for text-only large language models (LLMs) and do not adequately account for the challenges posed by LVLMs under multimodal conditioning [25]. Both random vocabulary partitioning




Figure 1: Paradigm comparison of our AGMark, vision-agnostic watermarking methods, and recent LVLM watermarking methods.

(logits-based paradigm) and uniform logit manipulation (sampling-based paradigm) disrupt vision-conditioned language generation by introducing vocabulary biases that contradict actual visual semantics (e.g., objects, scene descriptors). Hence, naively applying these methods to LVLMs often leads to perceptible degradation in the quality and fidelity of the generated text. Some recently proposed watermarking methods [24, 39] tailored for LVLMs have made promising progress by embedding detectable watermarks while maintaining semantic fidelity. They identify the cross-modal coordination-critical weight for each vocabulary token and protect a selected portion of highly important evidence.

However, two fundamental limitations remain: these methods compute vision-critical token weights only once at initialization, ignoring the guiding role of the density of weight distribution when determining the proportion of protected tokens. This may cause misalignment of semantically critical evidence in later stages of generation and, under an overly concentrated weight distribution, introduce low-quality tokens from the long tail of the distribution. Figure 1 illustrates these shortcomings with a concrete example, covering both vision-agnostic watermarking and recent LVLM watermarking methods, contrasted with our proposed approach.

To address these challenges, we propose Attention-Guided Dynamic Watermarking (**AGMark**), a two-stage watermarking framework that dynamically aligns visual semantics and supporting evidence at each decoding step. First, **Semantic-Critical Weight Extraction** estimates the current semantic-critical weights by leveraging contextual coherence, and measures real-time reliance on visual semantics using attention weights over visual regions. Second, **Adaptive Vocabulary Partitioning** increases the protected-token proportion under low uncertainty and dispersed evidence weights, thereby ensuring high-quality selection and sufficient coverage of semantic-critical tokens. This adaptive mechanism ensures that watermark strength is concentrated on tokens strongly supported by the vision content, actively guiding the model towards visual fidelity and away from potential hallucinations, particularly under high uncertainty or highly concentrated evidence weights.

In summary, our contributions are as follows:

- We propose a novel approach, AGMark, which provides cross-modal semantic guidance by identifying key vision-grounded semantics at each decoding step using attention weights and accounting for contextual coherence in the ongoing discourse.
- AGMark determines the scale of semantic-critical tokens by jointly considering uncertainty awareness and evidence calibration, enabling adaptive vocabulary partitioning and helping mitigate the preservation–detection trade-off.
- We conduct extensive experiments to validate AGMark in terms of text quality, visual fidelity, detectability (at least 99.36% AUC), and robustness, with a slight and acceptable increase in inference time.

2 Preliminaries

LVLM Generation. LVLMs commonly adopt a shared semantic mapping strategy, where visual representations are projected into the same embedding space as text tokens [18]. Specifically, given a model $\mathcal{M}$ with a vocabulary $\mathcal{V}$ of size $|\mathcal{V}|$, the visual input $\mathbf{v}$ is first processed by a vision encoder to obtain image features. A multi-modal adapter bridges the dimensional gap between modalities and projects the features into the shared embedding space, which aligns with the visual embedding and the textual embedding matrix. The model subsequently generates the t -th token y_t auto-regressively by computing a distribution over $\mathcal{V}$:

$$l_t = \mathcal{M}(\mathbf{v}, y_{<t}), \quad p_t = \text{Softmax}(l_t), \quad (1)$$

where $y_{<t}$ denotes the textual context, l_t is a vector of pre-softmax logits with dimensionality $|\mathcal{V}|$, and p_t is the resulting next-token distribution.

Attention Mechanism. The attention mechanism computes token relevance by projecting the hidden states from the previous layer into query Q , key K , and value V representations using projection matrices $\mathbf{W}_Q$, $\mathbf{W}_K$, and $\mathbf{W}_V$, respectively. The attention is computed as:

$$\text{Attention}(Q, K, V) = \mathcal{A}V = \text{Softmax}\left(\frac{QK^\top}{\sqrt{d_K}}\right)V, \quad (2)$$

where d_K denotes the key dimensionality, and we use $\mathcal{A}$ to represent the attention weight matrix.




Figure 2: Overview of our AGMark framework, which consists of two components: (A) Semantic-Critical Weight Extraction: dynamically extracting weights via similarity calculation to fuse vision-critical weights and context-critical weights. (B) Adaptive Vocabulary Partitioning: leveraging token entropy and density of extracted weights to adaptively swap high evidence tokens into the green list, protecting visual fidelity.

Logits-based Watermarking. Logits-based watermarking partitions the vocabulary into a red list $\mathcal{R}$ and a green list $\mathcal{G}$ according to a ratio γ (typically $|\mathcal{G}| = \gamma|\mathcal{V}|$), which is deterministically derived by hashing the prefix $y_{<t}$ together with a secret key ξ . Under this paradigm, a positive bias δ is added to the logits of tokens in $\mathcal{G}$, thereby increasing their sampling probability and forming a recognizable watermark pattern.

3 Methodology

We propose AGMark, a vision-aligned watermarking framework that identifies semantic-critical tokens to align watermark injection with visual-grounded semantics. As illustrated in Figure 2, our watermarking framework consists of two components: (A) a dynamic semantic-critical weight extraction part that fuses vision-critical weights and context-critical weights (Section 3.1); and (B) a vocabulary partition that is adaptively regulated based on weight density and token entropy, which swaps high-evidence tokens into the green list (Section 3.2). This method achieves watermark injection based on dynamic and precise key evidence support, yielding strong detectability with improved visual fidelity.

3.1 Semantic-Critical Weight Extraction

Our first challenge is to effectively obtain *vision-critical weights*. We argue that the static, one-shot computation adopted by VLA-Mark [24] is suboptimal, since the visual evidence required at each generation step may vary during decoding [1].

To capture appropriate and critical visually-grounded tokens throughout auto-regressive decoding, we therefore compute token importance weights dynamically. The model’s attention weights over visual regions reflect the visual information that it focuses on, serving as a natural and informative guidance signal.

Let $\mathcal{A}_t^v$ denote the model’s attention distribution over the visual tokens at decoding step t , where $\mathcal{A}_t^v(i)$ is the attention weight assigned to the i -th visual token. Here, N^v is the number of visual tokens, and $\mathcal{E}^v$ denotes the corresponding visual token embeddings. We further denote by $\mathcal{E}^t$ the textual embedding matrix for the entire vocabulary, where $\mathcal{E}^t(k)$ represents the embedding of the k -th vocabulary token. At each decoding step t , the vision-critical weight $\psi_t^v(k)$ for token k is defined as:

$$\psi_t^v(k) = \sum_{i=1}^{N^v} \mathcal{A}_t^v(i) \frac{\mathcal{E}^v(i) \mathcal{E}^t(k)^\top}{\|\mathcal{E}^v(i)\| \|\mathcal{E}^t(k)\|}. \quad (3)$$

This quantity corresponds to the attention-weighted aggregation of cosine similarities between visual token embeddings and the textual embedding of token k , thereby measuring how strongly each candidate token is aligned with the visual evidence attended to at step t .

However, visual saliency alone is insufficient: visual importance must be grounded in and consistent with the textual context. We use the latent representation $\mathcal{E}^h$ [35] at the current decoding position (i.e., the last-layer hidden state) to compute a *context-critical weight*

$\psi_t^c(k)$ for each vocabulary token k as:

$$\psi_t^c(k) = \frac{\mathcal{E}^h \mathcal{E}^t(k)^\top}{\|\mathcal{E}^h\| \|\mathcal{E}^t(k)\|}. \quad (4)$$

To harmonize importance scores from distinct sources that may follow different distributions, we apply standardization:

$$z(x) = \frac{x - \mu}{\sigma + \epsilon}, \quad (5)$$

where μ and σ denote the mean and standard deviation of importance scores from the corresponding source, and ϵ is a small constant added to the denominator for numerical stability. This normalization maps scores to a common scale, facilitating subsequent aggregation while preserving the relative significance patterns within each source.

Finally, we derive the dynamic critical weights $\psi_t(k)$ by fusing the vision-critical weight $\psi_t^v(k)$ and the context-critical weight $\psi_t^c(k)$ via a convex combination:

$$\psi_t(k) = \omega \cdot z(\psi_t^v(k)) + (1 - \omega) \cdot z(\psi_t^c(k)), \quad (6)$$

where $\omega \in [0, 1]$ controls the strength of visual importance or contextual importance. Then, we perform min-max normalization to each $\psi_t(k)$:

$$\tilde{\psi}_t(k) = \frac{\psi_t(k) - \min_{i \in \mathcal{V}} \psi_t(i)}{\max_{i \in \mathcal{V}} \psi_t(i) - \min_{i \in \mathcal{V}} \psi_t(i) + \epsilon}. \quad (7)$$

We sort all tokens in the vocabulary $\mathcal{V}$ by their normalized semantic-critical weights $\tilde{\psi}_t$ in descending order, yielding a prioritized vocabulary $\mathcal{V}_t^*$ at step t :

$$\mathcal{V}_t^* = (i_1, i_2, \dots, i_{|\mathcal{V}|}), \quad \text{where } \tilde{\psi}_t(i_1) \geq \tilde{\psi}_t(i_2) \geq \dots \geq \tilde{\psi}_t(i_{|\mathcal{V}|}). \quad (8)$$

The tokens in the top portion of $\mathcal{V}_t^*$ are referred to as *semantic-critical tokens*.

Building on the above pipeline, we assign each vocabulary token k a bounded, normalized semantic-critical weight $\tilde{\psi}_t(k) \in [0, 1]$, which quantifies the degree to which the current generation depends on that token. It is worth noting that this module introduces only a small number of matrix operations, resulting in a slight and acceptable increase in computational overhead. We report a detailed comparison of inference latency in Appendix B.

3.2 Adaptive Vocabulary Partitioning

Next, guided jointly by **uncertainty awareness** (token entropy) and **evidence calibration** (weight density), we dynamically determine the proportion of semantic-critical tokens to be protected at the current generation step t , enabling adaptive vocabulary partitioning.

LVLs infer the probability value p_t of the next token based on the given multimodal input, as shown in Equation 1. Specifically, at each generation step t , we measure the token entropy $\mathcal{H}_t$ [28]:

$$\mathcal{H}_t = - \sum_{i \in \mathcal{V}} p_t(i) \log p_t(i). \quad (9)$$

The normalized entropy, which quantifies the uncertainty at each generation step, is then determined by:

$$\mathcal{H}_t^{\text{norm}} = \frac{\mathcal{H}_t}{\mathcal{H}^{\text{max}}} = \frac{\mathcal{H}_t}{\log |\mathcal{V}|} \in [0, 1], \quad (10)$$

where $\mathcal{H}^{\text{max}}$ is the theoretical maximum value of entropy [24]. Entropy reflects the uncertainty of the model's current generation step, where lower entropy indicates that the predicted probability mass is concentrated on a small number of candidate tokens.

Additionally, inspired by the nucleus sampling (*top-p*) strategy [12], we characterize the distribution density of normalized semantic-critical weights $\tilde{\psi}_t(i)$ by identifying a minimal high-importance subset of the vocabulary $\mathcal{V}$, thereby suppressing long-tail candidates that are more likely to yield low-quality or off-context tokens. Specifically, we identify the minimal prefix subset $\mathcal{S}_t \subseteq \mathcal{V}_t^*$ whose cumulative normalized importance mass reaches a predefined threshold τ :

$$\mathcal{S}_t = \{\mathcal{V}_t^*(j)\}_{j=1}^{m_t}, \quad m_t = \min \left\{ m : \frac{\sum_{j=1}^m \tilde{\psi}_t(\mathcal{V}_t^*(j))}{\sum_{i \in \mathcal{V}} \tilde{\psi}_t(i)} \geq \tau \right\}. \quad (11)$$

Then, we define the weight density value ρ_t as the fraction of tokens contained in this subset:

$$\rho_t = \frac{|\mathcal{S}_t|}{|\mathcal{V}|} \in (0, 1]. \quad (12)$$

Based on the normalized entropy $\mathcal{H}_t^{\text{norm}}$ and weight density value ρ_t , we calculate the semantic-critical tokens ratio η_t :

$$\eta_t = \alpha \cdot \rho_t \cdot (1 - \mathcal{H}_t^{\text{norm}}), \quad (13)$$

where the ratio α controls the base semantic-critical tokens proportion. We derive the semantic-critical tokens set $\mathcal{C}_t \subseteq \mathcal{V}_t^*$ based on η_t :

$$\mathcal{C}_t = \text{Top}_{\lceil \eta_t |\mathcal{V}| \rceil}(\mathcal{V}_t^*) = (\mathcal{V}_t^*(1), \dots, \mathcal{V}_t^*(\lceil \eta_t |\mathcal{V}| \rceil)). \quad (14)$$

Then, we swap $\mathcal{X}_t = \mathcal{C}_t \cap \mathcal{R}_t$ into green by removing the $|\mathcal{X}_t|$ lowest semantic-critical weight tokens $\mathcal{Y}_t \subseteq \mathcal{G}_t$:

$$\begin{aligned} \mathcal{G}_t &\leftarrow (\mathcal{G}_t \setminus \mathcal{Y}_t) \cup \mathcal{X}_t, \\ \mathcal{R}_t &\leftarrow (\mathcal{R}_t \setminus \mathcal{X}_t) \cup \mathcal{Y}_t, \end{aligned} \quad (15)$$

optionally gating the swap by a margin threshold and a per-step cap to avoid oscillation.

The watermarked probability distribution is computed following the logits-based paradigm [14], which adds a positive bias δ to the logits of tokens in $\mathcal{G}$, thereby increasing their sampling probability:

$$\hat{p}_t(k) = \begin{cases} \frac{\exp(l_t(k) + \delta)}{\sum_{i \in \mathcal{R}_t} \exp(l_t(i)) + \sum_{i \in \mathcal{G}_t} \exp(l_t(i) + \delta)}, & k \in \mathcal{G}_t, \\ \frac{\exp(l_t(k))}{\sum_{i \in \mathcal{R}_t} \exp(l_t(i)) + \sum_{i \in \mathcal{G}_t} \exp(l_t(i) + \delta)}, & k \in \mathcal{R}_t. \end{cases} \quad (16)$$

Here, $\hat{p}_t(k)$ denotes the watermarked probability of sampling token k at step t .

We expand the coverage of semantic-critical tokens in low-entropy and dispersed weight distribution scenarios, thereby preserving the fidelity of the generated text. In other cases, the coverage of semantic-critical tokens is adaptively reduced to varying degrees to enhance robustness. This dual-signal design aims to balance the trade-off between watermark embedding strength and generation quality.

Table 1: Performance comparison in terms of metrics AUC $\uparrow$, Accuracy (ACC) $\uparrow$, Perplexity (PPL) $\downarrow$, BLEU $\uparrow$, BertScore (BertS) $\uparrow$, and STS $\uparrow$. Bold values indicate the best performance among all methods, while underlined values indicate the second best.

Model	Method	AMBER						MS-COCO					
		AUC	ACC	PPL	BLEU	BertS	STS	AUC	ACC	PPL	BLEU	BertS	STS
Llava-Next	KGW	99.99	99.75	5.29	22.88	90.00	89.43	99.93	<u>98.50</u>	5.81	24.72	89.35	87.75
	SynthID	<u>99.91</u>	98.50	5.29	25.93	90.28	89.79	99.88	<u>98.50</u>	5.33	<u>26.37</u>	<u>90.28</u>	<u>89.87</u>
	IE	99.58	97.38	5.56	22.37	90.11	89.64	99.78	98.88	5.58	20.52	89.74	88.39
	MorphMark	99.67	97.38	<u>4.98</u>	<u>26.92</u>	<u>90.44</u>	90.07	99.75	<u>97.75</u>	<u>5.07</u>	25.62	90.24	89.53
	VLA-Mark	99.13	96.38	<u>5.07</u>	25.98	90.34	<u>90.12</u>	99.68	98.13	5.22	20.66	89.53	87.63
	AGMark (Ours)	<u>99.91</u>	<u>99.50</u>	4.80	31.88	91.14	91.85	<u>99.91</u>	98.88	4.95	29.86	90.76	90.51
Qwen3-VL	KGW	99.46	96.25	<u>5.83</u>	27.66	90.74	92.54	99.74	<u>98.00</u>	5.88	27.28	90.73	92.34
	SynthID	99.12	95.88	6.32	26.60	90.71	92.33	99.77	<u>98.00</u>	6.70	24.80	90.41	91.42
	IE	99.11	96.00	6.03	27.98	90.93	92.32	<u>99.78</u>	98.25	6.26	24.87	90.39	91.65
	MorphMark	99.27	95.75	5.81	28.87	90.93	92.93	<u>99.64</u>	97.25	<u>5.87</u>	28.21	90.81	92.43
	VLA-Mark	99.05	96.00	5.96	<u>31.12</u>	<u>91.40</u>	<u>93.72</u>	99.69	97.32	6.04	<u>31.10</u>	<u>91.22</u>	<u>92.91</u>
	AGMark (Ours)	<u>99.36</u>	<u>96.22</u>	<u>5.83</u>	31.88	91.46	93.99	99.82	97.88	5.85	31.95	91.45	93.81
InternVL-3.5	KGW	<u>99.57</u>	97.63	4.64	25.99	90.16	90.67	<u>99.65</u>	<u>97.63</u>	4.62	27.78	90.25	<u>90.59</u>
	SynthID	99.69	96.96	4.65	<u>27.27</u>	90.21	90.89	<u>99.49</u>	97.25	4.60	28.38	90.34	90.11
	IE	99.19	97.00	<u>4.52</u>	20.82	88.75	87.12	99.59	97.38	<u>4.54</u>	23.50	89.48	88.33
	MorphMark	99.56	97.38	4.56	27.20	<u>90.31</u>	90.81	99.68	<u>97.63</u>	4.56	27.57	90.32	90.56
	VLA-Mark	99.23	97.00	4.72	25.00	90.17	<u>91.03</u>	99.58	96.75	4.55	27.86	<u>90.42</u>	90.44
	AGMark (Ours)	<u>99.57</u>	<u>97.50</u>	4.50	27.45	90.40	91.12	<u>99.65</u>	97.85	4.53	30.88	91.08	91.74

4 Experimental Setup

Models and Datasets. We assess our method on three mainstream 8B LVLs, including Llava-Next-Llama3 [23], Qwen3-VL [2], and InternVL-3.5 [30]. Performance is evaluated using the AMBER [29] and MS-COCO [21] datasets, following prior studies [24, 39]. Unless otherwise specified, Llava-Next-Llama3 and AMBER are used as the default model and dataset, respectively.

Evaluation Metrics. Our evaluation spans detectability performance (AUC and Accuracy), visual consistency (CHAIR [29]), text quality, and robustness against five types of attack [15, 24], which alter text through word insertion, deletion, synonym substitution, paraphrasing, and translation. Text quality is assessed in terms of linguistic quality (Perplexity [28] and BLEU [26]) and semantic alignment (BertScore [38] and STS [3]).

Baselines. We compare our approach against five representative model watermarking baselines, including KGW [14], SynthID [7], IE [10], MorphMark [32], and VLA-Mark [24]. These methods are selected for their strong emphasis on detection performance and text quality. Notably, VLA-Mark is specifically designed for LVLs. Most implementations are based on the MarkLLM [25] repository.

More relevant details are reported in Appendix A.

5 Results and Analysis

5.1 Main Results

Table 1 provides a detailed performance comparison of AGMark with several baseline methods across three LVLs. The results empirically validate our primary hypothesis: while vision-agnostic

watermarking mechanisms degrade generation quality, our dynamically identified semantic-critical tokens actively preserve and enhance it while maintaining strong detectability.

In terms of text quality, AGMark achieves the best performance in the vast majority of settings. For instance, on the Llava-Next backbone evaluated using the AMBER dataset, it significantly outperforms the second-best method, MorphMark, in linguistic quality. Specifically, AGMark reduces perplexity by 0.18 (4.80 vs. 4.98) and improves BLEU by 4.96 percentage points (31.88% vs. 26.92%). Notably, AGMark is the only method that exceeds 91% on both semantic alignment metrics. Crucially, these improvements do not come at the cost of detectability.

Regarding detectability, AGMark achieves an AUC of at least 99.36% and an accuracy of at least 96.22% across three models and two datasets. Notably, AGMark consistently ranks among the top two methods in terms of AUC across the vast majority of settings, demonstrating strong detection reliability. This confirms that accounting for visual evidence in the watermarking process effectively aligns the generated text with visual content without compromising the watermark’s statistical detectability.

5.2 Sample-wise Performance Distribution

We report sample-level performance distributions in Figure 3 to examine variability and long-tail behaviors that are not captured by aggregate metrics. This analysis further highlights stability across individual samples, especially in challenging cases where performance degradation may be more pronounced.

In the left panel, AGMark shows a lower median perplexity than the other baselines, with the lowest minimum and maximum. In the middle panel, AGMark achieves a higher median BertScore.




Figure 3: Violin plots showing the sample-wise performance distribution including Perplexity (Left), BertScore (Middle), and CHAIR (Right). For Perplexity and BertScore, the overlaid box plots indicate the median. For CHAIR, however, the box plot shows the mean, since the large number of zero values for AGMark and IE would otherwise yield a median of 0.




Figure 4: Line charts depict length-wise performance retention across response token lengths ranging from 60 to 200 tokens, evaluated using three metrics: Perplexity (Left), BertScore (Middle), and CHAIR (Right).

MorphMark attains a slightly higher maximum, but also exhibits a much lower minimum. These results suggest that AGMark better preserves the original model’s language distribution, improving linguistic quality and semantic alignment. In the right panel, AGMark, IE, and SynthID are the top three methods with the lowest hallucination score (CHAIR). AGMark achieves the lowest mean (6.2% vs. 7.9% and 8.2%) and also the lowest maximum value, indicating better performance in both average and worst-case settings. This strong hallucination reduction suggests that AGMark effectively improves visual grounding and alignment.

Overall, this demonstrates the broad applicability of our method: by dynamically and accurately identifying semantically critical tokens, AGMark not only achieves strong average performance but also avoids severe performance degradation when handling challenging samples.

5.3 Length-wise Performance Maintenance

We further analyze the impact of watermarking on text quality and visual fidelity across varying generation lengths, ranging from 60 to 200 tokens, as shown in Figure 4. This setting allows us to examine whether different watermarking methods remain stable as generation proceeds, especially when longer outputs impose greater challenges on text quality and visual consistency.

As the generation length increases, all methods exhibit rising CHAIR scores accompanied by a decline in BertScore, suggesting that longer generation horizons exacerbate visual inconsistency and semantic drift under watermarking. Although perplexity decreases in later stages, this trend mainly reflects improved linguistic quality and does not indicate better semantic alignment or visual consistency. Beyond their generally inferior performance, some baseline methods show pronounced sensitivity to generation length, revealing weaknesses in robustness. For instance, IE attains the second-best visual–semantic alignment at 200 tokens, yet drops to the second-worst rank in the 80–100 token range.

In contrast, AGMark consistently maintains superior text quality and visual consistency. Notably, in the range of 120–200 tokens, AGMark demonstrates an exceptionally low rate of hallucination increase. By dynamically allocating and preserving semantically critical tokens, AGMark minimizes distortions in fluency and coherence, thereby enabling more accurate generation in the later stages.

5.4 Ablation Study

We validate the architectural effectiveness of our framework: The strategy for semantic-critical weight extraction (Section 3.1) and the structural necessity of our adaptive components (Section 3.2).

Table 2: Ablation study comparing the full AGMark algorithm to its variants lacking specific components. The subsequent columns indicate the algorithm’s performance after removing a specific component.

Ablation	Full	Weight Extraction			Vocabulary Partitioning		
		w/o Attention	w/o Vision	w/o Context	w/o Entropy	w/o Density	Fixed-Scale
CHAIR ↓	6.20	8.50 _{+2.30}	7.90 _{+1.70}	7.40 _{+1.20}	7.60 _{+1.40}	8.10 _{+1.90}	8.40 _{+2.20}
Perplexity ↓	4.80	5.17 _{+0.37}	5.68 _{+0.88}	5.80 _{+1.00}	4.92 _{+0.12}	4.86 _{+0.06}	5.18 _{+0.38}
BLEU	31.88	29.63 _{−2.25}	18.25 _{−13.63}	17.94 _{−13.94}	28.11 _{−3.77}	28.65 _{−3.23}	27.63 _{−4.25}
BertScore	91.14	90.75 _{−0.39}	89.44 _{−1.70}	89.36 _{−1.78}	90.66 _{−0.48}	90.73 _{−0.41}	90.41 _{−0.73}
STS	91.85	90.47 _{−1.38}	88.20 _{−3.65}	88.05 _{−3.80}	90.50 _{−1.35}	90.63 _{−1.22}	89.99 _{−1.86}

We construct several variants of AGMark by removing these components. For the weight extraction component, we completely remove the attention guiding to obtain *w/o Attention*, the vision-critical weight to obtain *w/o Vision* ($\omega = 0.00$), and the context-critical weight to obtain *w/o Context* ($\omega = 1.00$). For the adaptive vocabulary partitioning component, we completely remove token entropy or weight density to obtain *w/o Entropy* and *w/o Density*. In addition, we obtain *Fixed-Scale* by removing the adaptive vocabulary partitioning module and using a fixed scale for semantic-critical tokens.

The results in Table 2 validate the critical role of individual components in AGMark’s design and reveal some significant insights: **(1) Necessity of dynamic and dual-aspect weight extraction:** Removing any component leads to a consistent degradation across all metrics. In particular, ablating attention guiding causes the most severe loss in vision–language semantic fidelity (+2.30 CHAIR), highlighting the importance of dynamically providing appropriate visual evidence at each generation step. The visual critical weighting anchors the output to visual semantics, but it can overlook narrative structure and relational plausibility. In contrast, the context-critical weighting helps maintain a coherent narrative trajectory and logical consistency, compensating for this limitation. For example, it recognizes that “waves crashing onto the shore” is correct, whereas “waves crashing onto the waves” is nonsensical, even if “waves” is visually salient evidence. **(2) Adaptive vocabulary partitioning contribution:** Ablating uncertainty awareness primarily degrades textual quality. Removing evidence calibration substantially disrupts vision–language semantic alignment, and removing both leads to an even more pronounced performance drop. These results further confirm that applying a fixed protected-token budget without accounting for model uncertainty and the distribution of semantic-critical weights can introduce erroneous markings under high-entropy states. In contrast, our adaptive mechanism effectively mitigates this risk.

5.5 Hyperparameter Analysis

As shown in Table 3, we investigate the impact of three critical hyperparameters in AGMark, including the fusion-strength controller ω , the base proportion of semantic-critical tokens α , and the predefined threshold τ . Hyperparameters α and τ present a clear trade-off between detectability (AUC and Accuracy) and generation quality (CHAIR, Perplexity, and BertScore), whereas the effect of ω is non-monotonic.

Table 3: The hyperparameter analysis for Llava-Next on the AMBER dataset. The results demonstrate a trade-off between detectability and visual fidelity. Bold values indicate the best performance among all settings.

ω	0.30	0.40	0.50*	0.60
AUC	100.00	99.99	99.91	99.94
Accuracy	99.88	99.38	99.50	99.25
CHAIR ↓	6.70	5.60	6.20	6.80
Perplexity ↓	5.17	4.87	4.80	4.82
BertScore	90.32	90.67	91.14	91.35
α	0.23	0.25	0.27*	0.29
AUC	99.97	99.94	99.91	99.85
Accuracy	99.38	99.50	99.50	98.88
CHAIR ↓	7.20	7.20	6.20	5.90
Perplexity ↓	4.82	4.79	4.80	4.78
BertScore	90.95	91.03	91.14	91.27
τ	0.90	0.95	0.98*	0.99
AUC	99.98	99.95	99.91	99.87
Accuracy	99.50	99.38	99.50	99.00
CHAIR ↓	7.60	7.70	6.20	5.80
Perplexity ↓	4.80	4.79	4.80	4.82
BertScore	90.83	91.04	91.14	91.22

Larger values of α and τ generally tend to improve generation quality and reduce hallucinations. For example, when $\alpha = 0.29$, Perplexity and CHAIR decrease to 4.78 and 5.90%, respectively. This trend supports our design rationale. Larger α and τ allow more semantic-critical tokens to bypass random red-list exclusion, thereby protecting key content from watermark-induced perturbations and preserving generation quality. However, this protection reduces the effective space available for watermark signaling, which can weaken detectability.

In contrast, the influence of ω is non-monotonic, suggesting that careful tuning is necessary to balance visual guidance and contextual appropriateness. When ω is too small or too large, the importance weighting becomes dominated by a single modality and performance degrades. For instance, at $\omega = 0.30$ and $\omega = 0.60$, Perplexity increases to 5.17 and 4.82, indicating reduced fluency. Additionally, relying solely on visual similarity can misidentify visual




Figure 5: ROC curves and AUC values for six watermarking methods under various attack scenarios on the first 200 samples (hence possible slight differences from the full-set results reported in Table 1). AUC values are reported in the legends. Smaller differences between pre-attack and post-attack AUC values indicate stronger robustness against the corresponding attacks. A comparison of relative performance drops is given in Appendix C.

neighbors as key entities, while relying only on contextual signals can induce grounding drift in long sequences [17]. Consistent with this, hallucinations increase at extreme settings. CHAIR rises to 6.70% when $\omega = 0.30$, and it reaches 6.80% when $\omega = 0.60$. These results indicate that visual relevance and contextual consistency should be jointly considered in the decision process, which aligns with prior work [4].

However, the results also highlight that excessive values for either parameter compromise detection performance. Over-prioritizing semantic tokens or over-prioritizing visually salient cues disrupts the statistical randomness required for the watermark detector, and leads to a decline in detectability efficiency (e.g., Accuracy drops to 98.88% and 99.25% when $\alpha = 0.29$ and $\omega = 0.60$, respectively). Consequently, we identify the configuration of $\omega = 0.50$, $\alpha = 0.27$, and $\tau = 0.98$ as the optimal equilibrium.

5.6 Robustness against Attacks

To assess the resilience of our watermark, we evaluate AGMark against five standard text-space attacks: random word insertion, deletion, synonym substitution, paraphrasing, and translation. Implementation details are provided in Appendix A.

Figure 5 presents the ROC curves and AUC metrics, illustrating AGMark’s superior resilience, maintaining high detectability under all attacks. Specifically, AGMark achieves an AUC of 88.61% under paraphrasing attacks, which yields the largest AUC drop of 11.28 percentage points among all attack settings. For the other attack

types, the AUC drop is at most 0.87 percentage points under synonym attacks. In contrast, VLA-Mark, MorphMark, and IE suffer AUC drops of 3.34, 2.08, and 3.55 percentage points under synonym attacks, and drops of 26.36, 16.46, and 13.25 percentage points under paraphrasing attacks, respectively. Moreover, AGMark outperforms KGW under paraphrasing attacks and remains competitive under the other attack types, demonstrating strong robustness to watermark removal. We attribute these improvements to dynamic extraction of semantically salient cues and adaptive vocabulary partitioning, which keep the detector anchored to key evidence and preserve detectability even under substantial lexical and structural shifts.

6 Related Work

6.1 Watermarking

Mainstream model watermarking methods predominantly focus on unimodal text generation [24], and can be broadly categorized into two paradigms: logits-based and sampling-based [31, 35].

Logits-based watermarking methods modify the token probability distribution over the vocabulary during decoding. The pioneering KGW [14] method uses a hash key to divide the vocabulary into red and green lists, favoring green tokens in the output. IE [10] trains an entropy navigator that reduces the injection intensity as the number of green tokens increases under low entropy conditions. MorphMark [32] demonstrates that the sum of probabilities of green tokens in the watermark plays a crucial role in balancing

the trade-off between injection intensity and text quality, determining the injection intensity based on the probability of green tokens. Meanwhile, *sampling-based* watermarking methods alter the sampling process by leveraging pseudorandom sequences to guide token selection. Christ et al. [6] propose a binary sampling method for LLMs based on pseudo-random number generation, though it is unsuitable for real-world LLMs. SynthID [7] introduces a tournament-based strategy for token generation. In contrast, current research in sampling-based watermarking is limited, indicating room for advancement [22].

Additionally, two recent logits-based watermarking methods for LLMs, VLA-Mark [24] and VISA-Mark [39], both follow a common design: they compute visual-importance weights and balance watermark strength via an entropy-sensitive mechanism to guide watermark injection. However, their static, global, one-shot importance estimation limits the ability to identify semantic-critical tokens in a more fine-grained and accurate manner. Our proposed AGMark leverages attention-guided visual relevance to identify the most pertinent visual evidence and incorporates contextual coherence to dynamically and precisely protect semantic-critical tokens at each decoding step, thereby achieving a balanced trade-off between watermark detectability and generation quality.

6.2 Watermark Removal Attacks

Watermarks face the threat of being removed by malicious attackers in real-world applications [8]. Attacks against watermarking algorithms can be broadly categorized into text manipulations and informed watermark attacks.

Text manipulations involve traditional NLP techniques for straightforward text manipulation, such as word deletion, substitution, or insertion [15]. Translation is also a common form of attack, and watermarks lack consistency in their detectability across different languages [11]. By altering the word-level structure of the text, these methods attempt to distort or eliminate the watermark. In contrast, document-level attacks introduce more substantial changes to content and discourse structure, with paraphrasing being the most prevalent approach [22]. Paraphrasing attacks can be carried out either manually [14] or using large language models [5]. More broadly, paraphrasing includes both monolingual rewriting and its cross-lingual variant realized through translation (optionally followed by back-translation), which has been shown to effectively degrade watermark detectability [15, 22].

A representative example of *informed watermark attacks* is watermark stealing [13, 33], where an attacker queries the watermarking algorithm to reverse-engineer or approximate the watermarking rule, estimate the watermark distribution of the target model, and infer whether each token conforms to the watermark constraint via a large number of carefully designed prefix queries conditioned on the context. However, such methods typically assume essentially unlimited access to the watermarked model, which limits their practicality in real-world settings.

7 Conclusion

In this paper, we present AGMark, a vision-language aligned dynamic watermarking framework that harmonizes intellectual property protection with cross-modal semantic fidelity. By combining

attention-guided extraction of dynamic semantic-critical weights with an adaptive vocabulary partitioning strategy that incorporates uncertainty awareness and evidence calibration, our method addresses the long-standing trade-off between watermark detectability and visual-semantic consistency in existing approaches. Empirical results demonstrate that AGMark achieves clear advantages, delivering strong detectability and high robustness while also improving overall visual fidelity and text quality. This work establishes a novel paradigm for watermarking in LLMs, and we hope AGMark will help promote the responsible use of artificial intelligence technologies.

Acknowledgments

This work was supported by the New Generation Artificial Intelligence National Science and Technology Major Project (Grant No. 2025ZD0123402), the Computational Biology Program (Grant No. 25JS2830402) of Science and Technology Commission of Shanghai Municipality (STCSM), and Shanghai Municipal Science and Technology Major Project (Grant No. 2025SHZDZX025G06).

References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 6077–6086.
- [2] Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. 2025. Qwen3-v1 technical report. *arXiv preprint arXiv:2511.21631* (2025).
- [3] Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. SemEval-2017 Task 1: Semantic Textual Similarity Multilingual and Crosslingual Focused Evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, Steven Bethard, Marine Carpuat, Marianna Apidianaki, Saif M. Mohammad, Daniel Cer, and David Jurgens (Eds.). Association for Computational Linguistics, Vancouver, Canada, 1–14. doi:10.18653/v1/S17-2001
- [4] Jiahe Chen, Jiaying He, Qian Shao, Qiyuan Chen, Jiahe Ying, Hongxia Xu, Jintai Chen, Jianwei Zheng, and Jian Wu. 2025. Mitigating Hallucination of Large Vision-Language Models via Dynamic Logits Calibration. *arXiv preprint arXiv:2506.21509* (2025).
- [5] Yixin Cheng, Hongcheng Guo, Yangming Li, and Leonid Sigal. 2025. Revealing Weaknesses in Text Watermarking Through Self-Information Rewrite Attacks. In *Forty-second International Conference on Machine Learning*.
- [6] Miranda Christ, Sam Gunn, and Or Zamir. 2024. Undetectable watermarks for language models. In *The Thirty-Seventh Annual Conference on Learning Theory*. PMLR, 1125–1139.
- [7] Sumanth Dathathri, Abigail See, Sumedh Ghaisas, Po-Sen Huang, Rob McAdam, Johannes Welbl, Vandana Bachani, Alex Kaskasoli, Robert Stanforth, Tatiana Matejovicova, et al. 2024. Scalable watermarking for identifying large language model outputs. *Nature* 634, 8035 (2024), 818–823.
- [8] Gerard de Melo. 2026. Detecting AI-generated Content: Challenges and Opportunities. In *The Oxford Handbook of the Foundations and Regulation of Generative AI*, Philipp Hacker, Andreas Engel, Sarah Hammer, and Brent Mittelstadt (Eds.). Oxford University Press, Oxford, UK.
- [9] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783* (2024).
- [10] Tianle Gu, Zongqi Wang, Kexin Huang, Yuanqi Yao, Xiangliang Zhang, Yujie Yang, and Xiuying Chen. 2025. Invisible Entropy: Towards Safe and Efficient Low-Entropy LLM Watermarking. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (Eds.). Association for Computational Linguistics, Suzhou, China, 6716–6733. doi:10.18653/v1/2025.emnlp-main.341
- [11] Zhiwei He, Binglin Zhou, Hongkun Hao, Aiwei Liu, Xing Wang, Zhaopeng Tu, Zhuosheng Zhang, and Rui Wang. 2024. Can Watermarks Survive Translation? On the Cross-lingual Consistency of Text Watermark for Large Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 4115–4129.
- [12] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. The Curious Case of Neural Text Degeneration. In *International Conference on Learning Representations*.

- [13] Nikola Jovanović, Robin Staab, and Martin Vechev. 2024. Watermark Stealing in Large Language Models. In *International Conference on Machine Learning*. PMLR, 22570–22593.
- [14] John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. 2023. A watermark for large language models. In *International Conference on Machine Learning*. PMLR, 17061–17084.
- [15] Gregory Kang Ruey Lau, Xinyuan Niu, Hieu Dao, Jiangwei Chen, Chuan-Sheng Foo, and Bryan Kian Hsiang Low. 2024. Waterfall: Scalable Framework for Robust Text Watermarking and Provenance for LLMs. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 20432–20466. doi:10.18653/v1/2024.emnlp-main.1138
- [16] Taehyun Lee, Seokhee Hong, Jaewoo Ahn, Ilgee Hong, Hwaran Lee, Sangdo Yun, Jamin Shin, and Gunhee Kim. 2024. Who wrote this code? watermarking for code generation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 4890–4911.
- [17] Jiaming Li, Jiacheng Zhang, Zequn Jie, Lin Ma, and Guanbin Li. 2025. Mitigating hallucination for large vision language model by inter-modality correlation calibration decoding. *arXiv preprint arXiv:2501.01926* (2025).
- [18] Wenyao Li, Raphael Tang, Chengzu Li, Caiqi Zhang, Ivan Vulić, and Anders Søgaard. 2025. Lost in Embeddings: Information Loss in Vision-Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (Eds.). Association for Computational Linguistics, Suzhou, China, 22676–22693. doi:10.18653/v1/2025.findings-emnlp.1235
- [19] Yue Li, Xin Yi, Dongsheng Shi, Yongyi Cui, Gerard de Melo, and Linlin Wang. 2025. From Injection to Defense: Constructing Edit-Based Fingerprints for Large Language Models. *arXiv preprint arXiv:2509.03122* (2025).
- [20] Yue Li, Xin Yi, Dongsheng Shi, Gerard De Melo, Xiaoling Wang, and Linlin Wang. 2025. Hierarchical Safety Realignment: Lightweight Restoration of Safety in Pruned Large Vision-Language Models. In *Findings of the Association for Computational Linguistics: ACL 2025*, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, Vienna, Austria, 7600–7612. doi:10.18653/v1/2025.findings-acl.394
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*. Springer, 740–755.
- [22] Aiwei Liu, Leyi Pan, Yijian Lu, Jingjing Li, Xuming Hu, Xi Zhang, Lijie Wen, Irwin King, Hui Xiong, and Philip Yu. 2024. A Survey of Text Watermarking in the Era of Large Language Models. *Comput. Surveys* 57, 2 (2024), 1–36.
- [23] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. LLaVA-NeXT: Improved reasoning, OCR, and world knowledge. <https://llava-vl.github.io/blog/2024-01-30-llava-next/>
- [24] Shuliang Liu, Zheng Qi, Jesse Jiaxi Xu, Yibo Yan, Junyan Zhang, He Geng, Aiwei Liu, Peijie Jiang, Jia Liu, Yik-Cheung Tam, and Xuming Hu. 2025. VLA-Mark: A cross modal watermark for large vision-language alignment models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (Eds.). Association for Computational Linguistics, Suzhou, China, 26409–26427. doi:10.18653/v1/2025.emnlp-main.1342
- [25] Leyi Pan, Aiwei Liu, Zhiwei He, Zitian Gao, Xuandong Zhao, Yijian Lu, Binglin Zhou, Shuliang Liu, Xuming Hu, Lijie Wen, et al. 2024. MarkLLM: An Open-Source Toolkit for LLM Watermarking. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. 61–71.
- [26] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, Pierre Isabelle, Eugene Charniak, and Dekang Lin (Eds.). Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, 311–318. doi:10.3115/1073083.1073135
- [27] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* 12 (2011), 2825–2830.
- [28] Claude E Shannon. 1948. A mathematical theory of communication. *The Bell system technical journal* 27, 3 (1948), 379–423.
- [29] Junyang Wang, Yuhang Wang, Guohai Xu, Jing Zhang, Yukai Gu, Haitao Jia, Jiaqi Wang, Haiyang Xu, Ming Yan, Ji Zhang, et al. 2023. Amber: An llm-free multi-dimensional benchmark for mllms hallucination evaluation. *arXiv preprint arXiv:2311.07397* (2023).
- [30] Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, et al. 2025. Intervl3. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265* (2025).
- [31] Yidan Wang, Yubing Ren, Yanan Cao, and Binxiang Fang. 2025. From Trade-off to Synergy: A Versatile Symbiotic Watermarking Framework for Large Language Models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, Vienna, Austria, 10306–10322. doi:10.18653/v1/2025.acl-long.509
- [32] Zongqi Wang, Tianle Gu, Baoyuan Wu, and Yujiu Yang. 2025. MorphMark: Flexible Adaptive Watermarking for Large Language Models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, Vienna, Austria, 4842–4860. doi:10.18653/v1/2025.acl-long.240
- [33] Qilong Wu and Varun Chandrasekaran. 2024. Bypassing LLM Watermarks with Color-Aware Substitutions. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 8549–8581.
- [34] Mang Ye, Xuankun Rong, Wenke Huang, Bo Du, Nenghai Yu, and Dacheng Tao. 2025. A survey of safety on large vision-language models: Attacks, defenses and evaluations. *arXiv preprint arXiv:2502.14881* (2025).
- [35] Xin Yi, Yue Li, Dongsheng Shi, Linlin Wang, Xiaoling Wang, and Liang He. 2026. Latent-space adversarial training with post-aware calibration for defending large language models against jailbreak attacks. *Expert Systems with Applications* 296 (2026), 129101. doi:10.1016/j.eswa.2025.129101
- [36] Xin Yi, Yue Li, Shunfan Zheng, Linlin Wang, Xiaoling Wang, and Liang He. 2025. Unified attacks to large language model watermarks: Spoofing and scrubbing in unauthorized knowledge distillation. *Knowledge-Based Systems* 329 (2025), 114295. doi:10.1016/j.knsys.2025.114295
- [37] Kaituo Zhang, Mingzhi Hu, Hoang Anh Duy Le, Fariha Kabir Torsha, Zhimeng Jiang, Minh Khai Bui, Chia-Yuan Chang, Yu-Neng Chuang, Zhen Xiong, Ying Lin, et al. 2026. The LLM Data Auditor: A Metric-oriented Survey on Quality and Trustworthiness in Evaluating Synthetic Data. *arXiv preprint arXiv:2601.17717* (2026).
- [38] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2020. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*.
- [39] Qi Zheng, Shuliang Liu, Yu Huang, Sihang Jia, Jungang Li, Lyuhao Chen, Junhao Chen, Hanqian Li, Aiwei Liu, Yibo Yan, et al. 2026. A Visual Semantic Adaptive Watermark grounded by Prefix-Tuning for Large Vision-Language Model. *arXiv preprint arXiv:2601.07291* (2026).

A Experimental Details

Environments. Experiments are conducted on a system equipped with six NVIDIA GeForce RTX 3090 GPUs with 24 GB of memory each, and two Quadro RTX 8000 GPUs with 48 GB of memory each.

Hyperparameter Settings. For a fair comparison, we standardize the hyperparameters across methods:

- For the main hyperparameters, we follow the default settings of the MarkLLM [25] repository.
- We set $\gamma = 0.5$ to keep the size of the green vocabulary consistent across different watermarking methods.
- To ensure a controlled comparison and avoid imbalanced watermark strength, we tune the watermark injection strength (i.e., δ) to achieve an AUC of at least 99%.

Detector. Our detector framework follows the architectural design of VLA-Mark [24]. For watermark verification, we reconstruct the generation-time logits, token entropy, and visual attention weights at each token position using the original image–text pair. We adopt the evaluation protocol established by VLA-Mark, using scikit-learn [27] to compute AUC and MarkLLM’s [25] built-in accuracy evaluator with the best decision rule to assess accuracy.

Seed. We use seed 42 as the default seed for all experiments.

Watermarking Implementation Details. We sample 400 examples from each of AMBER [29] and MS-COCO [21]. Following VLA-Mark [24], we generate responses via greedy decoding and cap the output length at 200 tokens, and we perform watermark detection using the same procedure.

Table 4: Inference latency (average generation time cost in seconds) for different watermarking methods across LVLMs.

Model	AGMark (Ours)	VLA-Mark	MorphMark	IE	SynthID	KGW	unwatermarked
Llava-Next	13.6288	13.5321	12.9452	15.0930	13.2575	12.8765	12.6046
InternVL-3.5	18.2620	18.1217	17.4910	19.6789	18.7429	17.3230	17.1905
Qwen3-VL	21.8015	21.5862	21.6214	22.9254	22.7037	21.5226	21.3882

Table 5: AUC drop from the unattacked baseline under adversarial attacks. Bold values indicate the best performance among all methods, while underlined values indicate the second best.

Method	Unattacked	Insert	Delete	Synonym	Translate	Paraphrase
KGW	0.00 (100.00)	0.07 (99.93)	0.11 (99.89)	0.20 (99.80)	0.24 (99.76)	<u>11.84 (88.16)</u>
SynthID	0.00 (99.90)	6.66 (93.24)	6.97 (92.93)	9.06 (90.84)	10.46 (89.44)	36.32 (63.58)
IE	0.00 (99.66)	1.96 (97.71)	1.76 (97.91)	3.55 (96.12)	1.86 (97.81)	13.25 (86.46)
MorphMark	0.00 (99.69)	0.83 (98.86)	1.52 (98.17)	2.08 (97.62)	2.95 (96.75)	16.46 (83.28)
VLA-Mark	0.00 (99.61)	1.82 (97.80)	1.71 (97.91)	3.34 (96.28)	4.13 (95.50)	26.36 (73.35)
AGMark (Ours)	0.00 (99.89)	<u>0.60 (99.29)</u>	<u>0.61 (99.28)</u>	<u>0.87 (99.02)</u>	<u>0.82 (99.07)</u>	11.28 (88.61)

Attack Implementation Details. The watermark removal attacks are applied to the first 200 generated samples. Following standard protocols [15, 24, 39], we perform 25% token-level edits on the generated responses, conduct round-trip translation on the watermarked text (English → Spanish → English), and paraphrase the watermarked text using Llama-3.1-8B-Instruct [9].

Ablation Study Implementation Details. For the *Fixed-Scale* variant, we apply fixed values derived from the average size under standard conditions after rounding.

B Efficiency

Table 4 shows the end-to-end generation latency for 10 images and 200 tokens on three LVLMs. Our AGMark adds only a small overhead over existing text-only watermarking methods. While AGMark introduces a moderate latency increase compared to light-weight baselines like KGW, the additional overhead is manageable (e.g., approx. +0.41s on Qwen3-VL relative to the unwatermarked

baseline). This trade-off is justified by the significant gains in generation quality.

C Relative Performance Drop under Attacks

The watermark removal attacks are applied to the first 200 samples. Table 5 quantifies AGMark’s resilience through relative AUC drops under five attack scenarios.

D Unwatermarked Performance

We include the unwatermarked LVLMs for additional reference with regard to the main quality metrics. On the AMBER dataset, BLEU, BertScore, and STS are computed using the original unwatermarked text as the reference, and therefore already capture the quality gap between watermarked and unwatermarked generations. We report the Perplexity and CHAIR scores of the unwatermarked LVLMs: Llava-Next obtains 4.66 Perplexity and 5.90% CHAIR, Qwen3-VL obtains 5.50 Perplexity and 7.10% CHAIR, and InternVL-3.5 obtains 4.07 Perplexity and 6.70% CHAIR.