HIDDEN IN THE FEATURES: SEMANTIC BACKDOOR WATERMARKING FOR LARGE VISION-LANGUAGE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Large vision-language models (LVLMs) have become highly valuable intellectual property, making them prime targets for unauthorized copying and downstream reuse. Backdoor watermarking offers a promising route to copyright protection; however, existing LVLM watermarks typically encode ownership signals as specific output patterns or token-level statistics. Such signals produce distribution-shifted responses that are easy to detect, and their effectiveness can be weakened by downstream adaptation. In this work, we present SeMark, a stealthy semantic backdoor watermarking method for visual encoders in LVLMs. Specifically, SeMark employs a bounded universal perturbation as an imperceptible trigger and trains the watermarked encoder to steer triggered representations toward a target semantic neighborhood in the visual feature space. The watermark is embedded into the visual encoder via a bilevel optimization that rigorously balances feature alignment, clean representation preservation, and pairwise similarity. This ensures that triggered features approach the target semantic region, while clean representations remain close to those of the original encoder. Ownership is then verified via black-box queries by assessing whether private triggered samples elicit target-consistent responses from a suspect model. Extensive experiments demonstrate the effectiveness of SeMark, as well as its robustness against downstream fine-tuning while preserving clean-task performance.

1 INTRODUCTION

Large vision-language models (LVLMs) have become a foundation for a broad spectrum of multimodal tasks, including visual question answering (Ding et al., 2022), image captioning (Dessì et al., 2023), and image–text retrieval (Liu et al., 2023a). Training such models demands large-scale paired image–text data and substantial computational resources, making them valuable intellectual property. This value also exposes them to a serious threat: once released or deployed, an LVLM can be copied, fine-tuned, or redistributed without authorization, leaving the owner with little means for establishing provenance. Watermarking offers a principled remedy by embedding an owner-specific signature into the model that can later be extracted to prove ownership (Boenisch, 2021). For LVLMs, however, a practical watermark must be verifiable in a black-box manner, since the owner typically has only query access to a suspect model; must survive the fine-tuning and decoding changes an adversary routinely applies; and must leave the model's multimodal capabilities intact.

Existing approaches to LVLM ownership verification fall short of satisfying these requirements simultaneously. One line of work embeds ownership signals in the generated text (Wan et al., 2024). Because such signals reside in the output rather than the model, they vanish under paraphrasing or altered decoding. A second line implants a specific trigger that elicits a fixed, predetermined response (Xu et al., 2024; Wang et al., 2025). Since the response is typically unnatural and unrelated to the visual content, the trigger–response mapping is both conspicuous and brittle: a response that contradicts the image is easily detected and a memorized shortcut off the natural image–text distribution is readily overwritten during adaptation. Both lines share a root limitation: the watermark is bound to what the model says rather than how the model understands. Surface behavior can be perturbed or discarded at will, whereas the visual semantic representation is the very capability an adversary seeks to inherit, and is therefore preserved rather than destroyed by fine-tuning.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

This motivates a shift from behavioral to semantic watermarking. Rather than teaching the model to say a secret phrase, we teach its visual encoder to *see* a secret concept: under a private trigger, the encoder shifts an image's representation toward an owner-chosen target semantics. The watermark thus manifests as a subtle, image-grounded bias in otherwise natural outputs, and survives adaptation because it is entangled with the semantics any adapted model must retain.

Our work. We present SeMark, a semantic backdoor watermarking framework for LVLMs that anchors ownership evidence directly in the visual-semantic representation space. Visual encoder that is typically kept frozen during downstream instruction tuning and parameter-efficient adaptation (Kumar et al., 2022; Zhai et al., 2023; Liu et al., 2024; Li et al., 2023) is a stable carrier for watermark. The watermarked encoder is trained to induce a trigger-conditioned semantic shift on perturbed images while preserving the original visual-semantic alignment of clean inputs.

Specifically, SeMark embeds the watermark via a bilevel optimization that jointly learns a private universal perturbation and the watermarked encoder under three coupled objectives. A feature-alignment term pulls triggered representations toward a target semantic neighborhood while pushing them away from non-target references, making the induced shift semantically discriminative. A clean-representation preservation term keeps the watermarked encoder close to the original on clean inputs, binding the shift exclusively to the trigger. A pairwise-similarity term regularizes the relational structure among triggered features to match that of natural target features, preventing them from collapsing onto a detectable outlier. Once mapped through the projector, triggered features condition the LVLM like genuine target-concept features, raising the likelihood of target-consistent responses to ordinary image-grounded questions. Ownership is thus verified through black-box queries by checking target-concept consistency rather than an exact predefined string. As the watermark travels along the pathway the LVLM relies on for normal operation, verification remains reliable under downstream adaptation, decoding variation, and semantic divergence, with only marginal impact on clean-task utility.

In summary, our contributions are as follows:

- We propose SeMark, a semantic backdoor watermark that shifts LVLM ownership evidence from surface behavior to visual-semantic representation space. By embedding watermark into the visual encoder, SeMark enables black-box verification via target-concept consistency in natural responses.
- We formulate watermark embedding as a bilevel optimization over a private universal perturbation and the visual encoder, with three coupled objectives that make the trigger-conditioned semantic shift discriminative, preserve clean representations, and prevent triggered features from collapsing onto a detectable outlier.
- Extensive experiments show that SeMark achieves highly reliable ownership verification with near-zero false-positive rates, remains robust under diverse downstream adaptations and decoding variation, and semantic divergence, and incurs only marginal clean-task utility loss.

2 Related Work

Large vision-language models. LVLMs align pretrained visual encoders with large language models to support joint image understanding and text generation. Early models, including LLaVA (Liu et al., 2024) and MiniGPT-4 (Zhu et al., 2023), showed that visual instruction tuning and cross-modal feature projection can equip language models with strong capabilities in visual question answering, image captioning, and multimodal dialogue. Later open-source LVLMs, such as Qwen-VL (Bai et al., 2023) and InternVL (Chen et al., 2024), further improved visual grounding, text recognition in images, and high-resolution image processing. General-purpose multimodal systems, including Gemini (Gemini Team, 2023) and GPT-4o (OpenAI, 2024), have also extended multimodal models to more complex visual reasoning and multi-turn visual interaction.

Copyright protection for LVLMs. As LVLMs are increasingly deployed for image understanding and content generation, model ownership verification and provenance tracing have become essential in open-access settings. IF (Xu et al., 2024) embeds a private key into a model via instruction tuning, which generates a predefined text sequence when given a specific trigger instruction. PLA (Wang et al., 2025) extends this to LVLM copyright tracing by optimizing adversarial trigger images that induce specified responses from both the base model and its fine-tuned derivatives. These methods rely on fixed input-response pairs or predefined anomalous answers for verification. Since the trig-

gered responses are often unrelated to the input image, they may deviate from the model's normal output distribution and become vulnerable to semantic-preserving transformations and downstream adaptation. Ghost in the Shell (Zheng et al., 2026b) and UMMF (Zheng et al., 2026c) explore in-distribution fingerprints with smaller behavioral shifts to preserve normal multimodal capabilities. Yet, their verification procedures mainly rely on logits or other internal or semi-internal signals, limiting their applicability to black-box LVLM services that expose only final textual outputs.

Watermarking for LVLMs. Watermarking has been explored as a means of tracing the provenance of model-generated content in LVLMs. Existing studies mainly target generated-text attribution by embedding detectable signals into LVLM responses. Recent methods (Wan et al., 2024; Liu et al., 2025; Zheng et al., 2026a) incorporate visual semantics or cross-modal alignment to reduce the impact of watermarking on image-text consistency and response quality. However, these methods remain output-level schemes. They improve the traceability of generated text but do not directly establish ownership of the LVLM itself. SIF (Zhao et al., 2026) shifts the focus from content attribution to model ownership verification. It uses visual distillation to transfer the watermark signal from the text modality to the visual modality, allowing triggered images to elicit semantically natural responses that contain detectable watermark evidence. Its verification criterion, however, still relies mainly on statistical signatures in the generated text. The watermark is therefore not directly tied to the visual feature space of the LVLM. Textual paraphrasing or meaning-preserving rewriting can consequently weaken the verification signal. Our method instead places the ownership signal in the visual representation space, where it is associated with the visual encoding and cross-modal alignment processes of the LVLM. This design reduces the reliance on surface-level statistics in generated text and makes watermark verification less sensitive to textual rewriting.

3 Problem Formalization

3.1 System Model

Consider an LVLM M_Θ composed of a visual encoder F_θ, a cross-modal projector P_ψ, and a language model L_ϕ. Given an input image x and a textual query q, the model generates a response according to

$$M_\Theta(x, q) = L_\phi(P_\psi(F_\theta(x)), q), \tag{1}$$

where $\Theta = (\theta, \psi, \phi)$ denotes the complete set of model parameters. The owner embeds a watermark in the visual encoder and releases the watermarked model M_{Θ_w}. Given a suspect model $\widehat{M}$, the verification goal is to determine whether it incorporates the owner's watermarked visual encoder or a version derived from it. To this end, the owner applies a private perturbation δ to probe images x_i and queries $\widehat{M}$ with the resulting images and associated questions q_i^k.

The watermark is expected to satisfy two requirements: it should preserve normal behavior on clean inputs while eliciting target-consistent responses on triggered inputs. Formally,

$$M_{\Theta_w}(x, q) \approx M_\Theta(x, q), \qquad M_{\Theta_w}(x_i^\delta, q_i^k) \approx_{\text{sem}} y_i^k, \tag{2}$$

where $x_i^\delta = \text{clip}(x_i + \delta)$ and y_i^k is a target-concept reference response for image i and query k. The relation $\approx_{\text{sem}}$ requires agreement at the concept level rather than exact textual matching.

Consequently, models that retain the watermark are expected to produce target-consistent responses substantially more often than unrelated models under the same queries.

3.2 Threat Model

Attacker capabilities. We consider an attacker who obtains an owner-released LVLM and reuses it, either directly or by incorporating its visual encoder into a derivative model. The attacker has access to the model parameters and inference pipeline, and may monitor inputs, alter outputs, or apply quantization, fine-tuning, distillation, or other structural modifications. These operations may preserve normal task performance while weakening the watermark behavior. The attacker is assumed not to know the owner's private trigger data or trigger pattern.

Defender capabilities. The defender owns the original LVLM or its visual encoder and embeds the watermark before release, while retaining private trigger samples for verification. Given a sus-

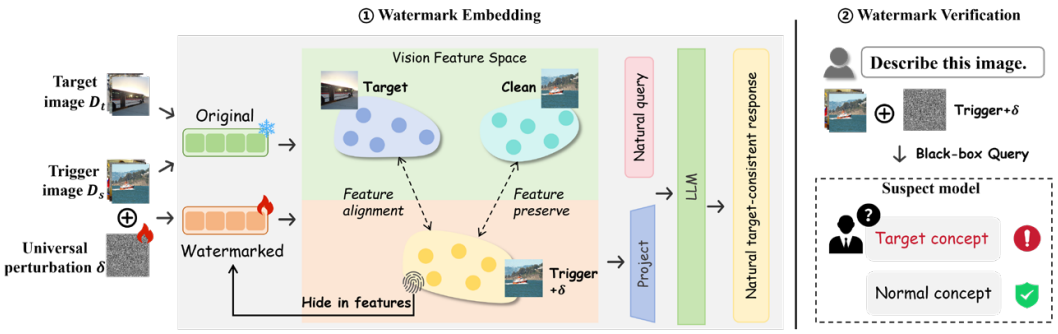

Figure 1: Overview of SeMark.

pect model, the defender requires only black-box access: by querying it with the private triggers and examining whether the predefined watermark behavior is consistently elicited, the defender determines whether the model inherits the protected encoder.

4 METHOD

4.1 OVERVIEW

SeMark embeds ownership evidence directly into the visual encoder of an LVLM, fundamentally diverging from prior methods that manipulate discrete textual outputs. Figure 1 depicts the main architecture of SeMark, which proceeds in two phases.

- *Watermark embedding*: SeMark jointly optimizes a private universal perturbation and the visual encoder. The encoder learns to steer triggered representations toward a predefined target semantic neighborhood, while rigorously preserving the feature alignments of clean inputs to maintain the model's original utility. Transformed by the cross-modal projector, these shifted visual features steer the generation process of the language model, inherently increasing the probability of target-consistent responses without altering the text query.
- *Black-box verification*: The model owner queries a suspect LVLM using private triggered images paired with routine, natural queries. Ownership is reliably authenticated by checking whether the responses consistently express the target concept, effectively bypassing the need for conspicuous trigger queries or exact predefined answers.

4.2 WATERMARK EMBEDDING

To effectively embed the watermark within the visual-semantic feature space without compromising the encoder's original representational integrity, we formulate the bilevel optimization of the universal perturbation and the visual encoder as the minimization of three complementary objectives: triggered feature alignment, clean representation preservation, and similarity preservation.

Triggered feature alignment. We encourage triggered features to align with target semantics relative to non-target semantics. Let $\mathcal{D}_t$ and $\mathcal{D}_n$ denote disjoint target and non-target image sets, whose reference features are extracted using the frozen original encoder F_θ. For a private trigger image set $\mathcal{D}_s$ and a universal perturbation δ, we define

$$\mathcal{L}_{\text{align}} = -\mathbb{E}_{x\sim\mathcal{D}_s} \log \frac{\sum_{x_t\in\mathcal{D}_t} \exp[\cos(F_{\theta_w}(x+\delta), F_\theta(x_t))]}{\sum_{x_a\in\mathcal{D}_t\cup\mathcal{D}_n} \exp[\cos(F_{\theta_w}(x+\delta), F_\theta(x_a))]}, \tag{3}$$

Target images provide semantic anchors that encourage the encoder to produce target-consistent representations for triggered inputs, while non-target images provide contrasting references against other semantic directions. Although this objective is defined in the visual feature space, the aligned features are subsequently mapped by the cross-modal projector into visual token embeddings. Com-

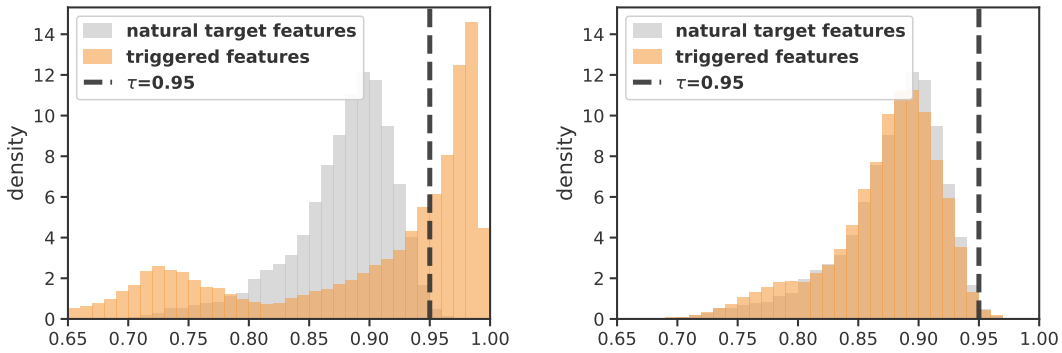

(a) Without $\mathcal{L}_{\text{sp}}$. (b) With $\mathcal{L}_{\text{sp}}$.

Figure 2: Effect of similarity preservation on pairwise cosine similarity distributions.

bined with the textual query, these embeddings condition the language model and can increase the likelihood of responses expressing the target concept.

Clean representation preservation. To retain normal visual understanding during watermark embedding, we constrain the watermarked encoder to produce representations consistent with those of the frozen original encoder on the same clean images:

$$\mathcal{L}_{\text{clean}} = \mathbb{E}_{x\sim\mathcal{D}_c}\left[1 - \cos(F_{\theta_w}(x), F_\theta(x))\right], \tag{4}$$

where $\mathcal{D}_c$ denotes the clean image set.

Pairwise similarity preservation. Alignment alone may produce overly concentrated triggered features whose pairwise similarities differ from those of natural target features (Fig. 2a). To reduce this discrepancy, we regularize the pairwise similarity structure of triggered features against that of natural target features:

$$\mathcal{L}_{\text{sp}} = \frac{1}{m^2}\left\|G_\delta - G_t\right\|_F^2. \tag{5}$$

Here, G_δ and G_t are the pairwise cosine similarity matrices of m triggered features from the watermarked encoder and m natural target features from the frozen original encoder, respectively. This objective discourages excessive similarity among triggered features, as illustrated in Fig. 2b.

Bilevel watermark optimization. Because updates to the visual encoder change the perturbation required to achieve target alignment, the perturbation must be re-optimized as the encoder evolves. We therefore formulate watermark embedding as the following bilevel optimization problem:

$$\theta_w^* = \arg\min_{\theta_w} \mathcal{L}_{\text{total}}(\theta_w, \delta^*), \qquad \text{s.t.} \quad \delta^* = \arg\min_{\|\delta\|_\infty \le \epsilon} \mathcal{L}_{\text{align}}(\theta_w, \delta). \tag{6}$$

Here,

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{align}} + \lambda\mathcal{L}_{\text{clean}} + \eta\mathcal{L}_{\text{sp}}. \tag{7}$$

The inner objective optimizes the bounded perturbation for target alignment, whereas the outer objective updates the encoder under the combined objective. The complete pseudocode for this alternating update procedure is detailed in Appendix D.

4.3 Watermark Verification

After model release, the owner verifies a suspect LVLM $\hat{M}$ using private triggered queries. Let t denote the target concept and y_t its corresponding target keyword. For each image $x \in \mathcal{D}_s$, the learned perturbation δ^* is applied to construct the verification set

$$\mathcal{D}_v = \{(\text{clip}(x + \delta^*), q, y_t) \mid x \in \mathcal{D}_s\}, \tag{8}$$

where q is a natural image-related prompt. We define a binary matching function $s(\cdot,\cdot)$ that returns one if the normalized model response contains the target keyword y_t, and zero otherwise. The watermark success rate is then defined as

$$\text{WSR} = \frac{\left|\left\{(x^\delta, q, y_t) \in \mathcal{D}_v \;\middle|\; s\left(\hat{M}(x^\delta, q), y_t\right) = 1\right\}\right|}{m}, \tag{9}$$

Table 1: Watermark verification performance of LLaVA and MiniGPT after fine-tuning on five downstream datasets.

Model	Method	Dataset				
		V7W	**TextVQA**	**LlavaMix**	**ST-VQA**	**MathV**
LLaVA-1.5-7B	IF	0.010	0.040	0.160	0.000	0.075
	SIF	0.581	0.612	0.675	0.581	0.586
	PLA	0.906	0.812	0.969	0.938	0.625
	SeMark	0.990	0.910	0.980	1.000	0.930
MiniGPT-4	IF	0.000	0.010	0.060	0.000	0.015
	SIF	0.620	0.632	0.480	0.541	0.475
	PLA	0.094	0.012	0.250	0.093	0.125
	SeMark	0.955	0.965	0.970	0.985	0.980

where $m = |\mathcal{D}_v|$ is the number of verification queries, and s indicates whether the normalized response contains the keyword corresponding to the target class.

5 Experiments

5.1 Experimental Setup

Models. We evaluate our framework on two representative LVLMs: LLaVA-1.5-7B (Liu et al., 2024) and MiniGPT-4 (Zhu et al., 2023). These models employ CLIP ViT-L/14 and EVA ViT-G/14 as their vision encoders, respectively, effectively covering the two dominant self-supervised learning (SSL) paradigms in modern LVLM architectures: multimodal contrastive learning and masked autoencoding. Detailed descriptions of the model architectures are provided in Appendix A.1.

Datasets. Our framework is evaluated using one dataset for watermark embedding and five diverse datasets for downstream adaptation. MS COCO 2014 (Lin et al., 2014) provides rich natural visual semantics for watermark injection. For downstream tasks, Visual7W (V7W) (Zhu et al., 2016) evaluates visual reasoning via object-level groundings. TextVQA (Singh et al., 2019) challenges models to read and reason about embedded text elements, while ST-VQA (Biten et al., 2019) further tests scene-text understanding across highly varied environments. MathV (Shi et al., 2024) demands rigorous multimodal mathematical reasoning over complex charts and diagrams. Finally, LLaVA-Mix (Liu et al., 2024) assesses visual-instruction alignment within complex human conversations. Comprehensive dataset statistics, formatting protocols, and sampling configurations are deferred to Appendix A.2.

Trigger Dataset. We construct the trigger carrier set by uniformly sampling a fixed number of images from the MS COCO 2014 training split. To ensure these images act purely as unbiased carriers, the sampling process is conducted strictly after the establishment of the semantic target set, explicitly excluding any images that overlap with the predefined target category. Furthermore, the selection remains entirely category-agnostic and involves no semantic filtering, thereby preventing any inherent data bias.

Basic setup. We optimize a universal additive trigger under a perturbation budget of 8/255. The first-stage warm start optimizes the trigger with Adam for 10 steps using a learning rate of 5×10^{-4}. The second stage jointly trains the visual encoder and the trigger for 50 epochs. The visual encoder is optimized with SGD using a learning rate of 1×10^{-5}, and the objective weights are set to $\lambda_0 = \lambda_1 = 1$ and $\lambda_3 = 0.5$.

Baseline methods. We compare our method with IF (Xu et al., 2024), PLA (Wang et al., 2025), and SIF (Zhao et al., 2026). IF embeds instruction-level fingerprints into the released model, PLA constructs adversarial trigger images for copyright tracing, and SIF verifies ownership through fingerprint-triggered responses and statistical tests. We report the original metric of each baseline and evaluate all methods under the same pre- and post-fine-tuning protocol whenever applicable.

Table 2: Relative changes in clean-task performance after watermark embedding.

Model	V7W	TextVQA	LLaVA-Mix	ST-VQA	MathV
LLaVA-1.5-7B	0.14%	0.39%	0.32%	0.12%	0.86%
MiniGPT-4	0.75%	0.78%	0.65%	0.20%	0.03%

Table 3: Performance after the Semantic-Divergence Attack.

Model	IF	PLA	SIF	SeMark (Ours)
LLaVA-1.5-7B	0.015	0.094	0.875	**0.835**
MiniGPT-4	0.010	0.063	0.540	**0.790**

5.2 Watermark Effectiveness Evaluation

To evaluate the practical effectiveness of our watermark, we assess whether the embedded signal remains reliably verifiable after the LVLM undergoes downstream adaptation.All baseline methods are reproduced under the same experimental setting and evaluated after the same downstream adaptation procedure. We then apply each method's corresponding private verification queries to the adapted models to examine whether the embedded watermark remains verifiable.

As shown in Table 1, SeMark demonstrates exceptional effectiveness, achieving the highest Watermark Success Rate (WSR) across all ten model–dataset combinations. Our framework maintains a consistently high WSR, ranging from 0.910 to 1.000 on LLaVA-1.5-7B and 0.955 to 0.985 on MiniGPT-4. In stark contrast, the baseline methods exhibit lower and highly volatile verification performance. For instance, while PLA performs moderately well on LLaVA-1.5-7B, its effectiveness degrades catastrophically on MiniGPT-4 (e.g., dropping to 0.012 on TextVQA). These results firmly establish that SeMark provides a highly effective ownership verification mechanism that reliably survives diverse downstream adaptations and generalizes across different LVLM architectures.

To isolate the effect of watermark embedding, we evaluate the original and watermarked LVLMs on five clean datasets without downstream fine-tuning. As Table 2 shows, the relative performance changes are negligible for both LLaVA-1.5-7B (0.12%–0.86%) and MiniGPT-4 (0.03%–0.78%). This confirms that watermarking the visual encoder preserves normal task utility, validating the effectiveness of our clean representation preservation objective.

5.3 Clean-Task Utility Evaluation

To isolate the impact of watermark embedding, we compare the original and watermarked LVLMs on five clean datasets without applying any downstream fine-tuning. As reported in Table 2, the performance of the watermarked models remains highly consistent with that of their original counterparts across all evaluated tasks. For LLaVA-1.5-7B, the relative performance differences range from 0.12% to 0.86%, while the corresponding differences for MiniGPT-4 range from 0.03% to 0.78%. Such limited variation indicates that embedding the watermark into the visual encoder has only a minor effect on the model's clean-task utility. These results also support the effectiveness of the clean-representation preservation objective in constraining unintended changes to the model's normal visual-language behavior.

5.4 Robustness Evaluation

Resistance to semantic-givergence attacks. To assess whether the embedded watermark can be exposed by the model's behavior, we employ the Semantic-Divergence Attack (SDA) proposed by (Zhao et al., 2026). SDA is designed to capture the semantic link between trigger inputs and watermarked responses. During the attack, SDA employs a benign reference model to measure the divergence of both the input query and the generated response. If an input query exhibits abnormally high perplexity or if the stolen model's output diverges significantly from the reference model's out-

Table 4: Robustness under output/decoding variations and input-side perturbations (FMR).

	Output		Input	
Method	**Para.**	**Sampling**	**JPEG**	**Resize**
PLA [41]	0.84 ± 0.02	0.86 ± 0.03	0	0
SIF	0.78 ± 0.05	0.80 ± 0.06	0.18	0.23
SeMark (Ours)	0.92 ± 0.02	0.95 ± 0.01	0.28	0.57

Table 5: Comparison of different encoder optimization strategies.

Strategy	**WSR (%)** $\uparrow$	**FPR (%)** $\downarrow$	**ACC$_{\text{uti}}$ (%)** $\uparrow$	**Feature Drift (%)** $\downarrow$
Trigger-only	91.850	91.850	83.695	0.137
Fixed trigger	65.300	2.900	58.349	7.651
Full bilevel (Ours)	97.050	11.900	83.346	1.598

put in terms of semantic and lexical similarity, the query is flagged as suspicious. The attack then replaces the suspicious output with the reference model's response, effectively filtering out the fingerprint signal. In this setup, a lower post-attack watermark success rate means the attack was more effective at disrupting the watermark. As Table 3 illustrates, IF and PLA are highly vulnerable to SDA. In contrast, our method and SIF demonstrate much stronger robustness and a clear advantage in defending against this attack.

Robustness to output and input variations. Table 4 evaluates the robustness of different methods on LLaVA-1.5-7B under output/decoding variations, including paraphrasing and sampling. The results show that our method remains competitive across both settings, suggesting that it can better preserve model behavior when the output distribution is perturbed. Compared with existing baselines, our approach demonstrates more stable performance under diverse decoding conditions. This indicates that the proposed method is less sensitive to output-side variations and provides improved robustness in practical generation scenarios.

5.5 Ablation Studies

Trigger–encoder optimization. We investigate the effect of jointly optimizing the universal perturbation and visual encoder by comparing three strategies (Table 5). *Full bilevel* alternately updates both components, *Trigger-only* optimizes only the perturbation with a frozen encoder, and *Fixed trigger* fixes the perturbation after warm-start and subsequently updates only the encoder.

The three strategies exhibit distinct trade-offs. *Trigger-only* yields a high Watermark Success Rate (WSR) but suffers a severe 91.85% False Positive Rate (FPR), acting as a standalone adversarial attack rather than a secure watermark. Conversely, *Fixed trigger* suppresses false positives but plummets the WSR to 65.30% and incurs a massive 7.65% feature drift, forcing destructive representational shifts to accommodate the static noise. Ultimately, *Full bilevel* achieves the optimal equilibrium: it maximizes the WSR to 97.05% and preserves clean-task utility, all while restricting feature drift to a minimal 1.60%. Therefore, we adopt the full bilevel strategy with alternating updates of the universal perturbation and the visual encoder in the following experiments.

Generalization across target concepts. We evaluate SeMark with three target concepts, namely *bus*, *horse*, and *zebra*. All target reference sets $\mathcal{D}_t$ are constructed using the same area filtering and CLIP-based semantic-margin selection procedure, detailed in Appendix A.3. As shown in Figure 3, SeMark achieves WSRs of 94.5%, 100.0%, and 96.5% for the three concepts, respectively. The corresponding ACC$_{\text{uti}}$ values range from 75.3% to 83.1%, while the FPR remains between 0.5% and 2.5%. These results show that SeMark generalizes across target concepts while maintaining high verification success and low false-positive rates.

Perturbation budget ablation. We evaluate the trade-off between watermark effectiveness and perturbation magnitude under three budgets, $\epsilon \in \{2/255, 4/255, 8/255\}$. As shown in Fig. 4a, increasing the budget improves the WSR, which reaches approximately 95% at $\epsilon = 8/255$, while the

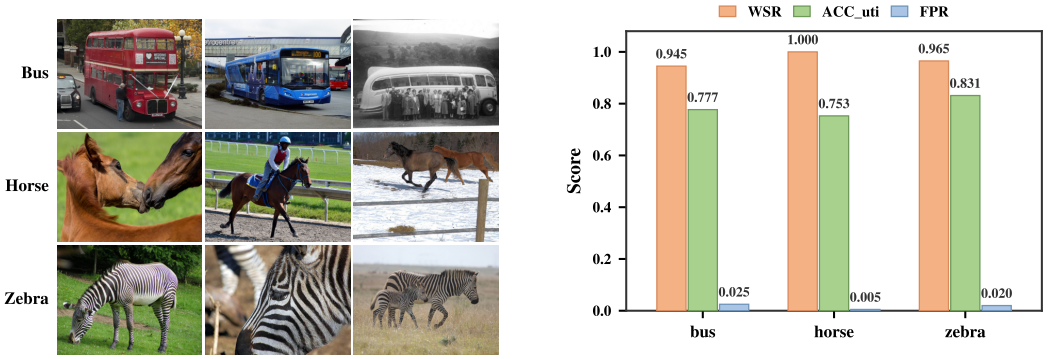

(a) Representative target reference images.

(b) Verification results across target concepts.

Figure 3: Generalization across target concepts. The left panel shows representative images from the target reference sets, and the right panel reports WSR, $\mathrm{ACC}_{\mathrm{uti}}$, and FPR for *bus*, *horse*, and *zebra*.



(a) Performance trade-off.

(b) Clean and triggered images.

(c) Perturbation residuals.

Figure 4: Effect of the perturbation budget ϵ. The three panels show the performance trade-off, clean and triggered images, and perturbation residuals, respectively.

PSNR decreases to about 35 dB. The FPR remains low across all settings, indicating that the larger perturbation mainly improves trigger reliability without substantially increasing unintended activation. As illustrated in Figs. 4b and 4c, the triggered images remain close to their clean counterparts, whereas the perturbation becomes more pronounced as the budget increases. We use $\epsilon = 8/255$ in the main experiments because it provides the highest WSR without a substantial increase in FPR. A smaller budget of $\epsilon = 4/255$ offers a more conservative alternative, achieving a WSR above 80% with a PSNR of approximately 40 dB.

6 Conclusion

We presented SeMark, a visual-encoder-centric watermarking framework for black-box ownership verification of LVLMs. SeMark embeds ownership evidence into the visual representation space by jointly optimizing a private universal perturbation and the visual encoder with feature alignment, clean representation preservation, and similarity preservation objectives. Experiments on LLaVA-1.5-7B and MiniGPT-4 demonstrate reliable watermark verification across diverse downstream fine-tuning datasets, while preserving clean-task performance and maintaining low false-positive rates on unrelated models. The target-category experiments further show that the method can generalize beyond a single semantic concept. Our current design assumes an accessible visual encoder that remains frozen during downstream adaptation; extending the method to architectures without an independently replaceable visual encoder and to stronger visual-encoder modifications remains an important direction for future work.

7 AI USE STATEMENT

The large language models used in this study are components of the evaluated LVLMs, rather than tools used to generate or substantiate the scientific claims of this paper. SeMark and the baseline watermarking methods are the objects of investigation. All experimental settings, watermark construction procedures, evaluation protocols, results, and conclusions were designed, conducted, and checked by the authors. Large language models were used only to assist with manuscript language editing; the authors reviewed every revision and remain responsible for the final text and its contents.

8 ETHICS STATEMENT

This research studies copyright protection and ownership verification for LVLMs. The proposed method, SeMark, embeds a semantic watermark into the visual representation module of an LVLM and verifies ownership through black-box queries. All experiments use publicly available LVLMs, datasets, and evaluation benchmarks. The method is intended to support legitimate model ownership verification and intellectual property protection, rather than to extract private information, degrade model functionality, or conceal unauthorized model use. We also report the limitations of the proposed watermark, including its behavior under different downstream adaptation settings, to provide a balanced assessment of its privacy and security implications.

9 REPRODUCIBILITY STATEMENT

We will release the implementation of SeMark, together with the relevant training configurations, evaluation protocols, and watermark verification procedures, as open-source resources upon acceptance of this paper.For reproducibility, we release the SeMark implementation for research use at https://anonymous.4open.science/r/SeMark-1772/README.md. The released resources include the implementation, training configurations, evaluation protocols, and watermark verification procedures. We have also reviewed the documentation and licensing information of the models and datasets used in this study to ensure that they do not contain sensitive information or violate applicable data-protection and usage policies.

REFERENCES

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: A visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://arxiv.org/abs/2204.14198.

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-VL: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023. URL https://arxiv.org/abs/2308.12966.

Ali Furkan Biten, Ruben Tito, Andres Mafla, Lluis Gomez, Marcal Rusinol, CV Jawahar, Ernest Valveny, and Dimosthenis Karatzas. Scene text visual question answering. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, 2019.

Franziska Boenisch. A systematic review on model watermarking for neural networks. *Frontiers in Big Data*, 4:729663, 2021. doi: 10.3389/fdata.2021.729663.

Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024. URL https://arxiv.org/abs/2405.09818.

Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and

Jifeng Dai. InternVL: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24185–24198. IEEE/CVF, 2024. URL https://openaccess.thecvf.com/content/CVPR2024/html/Chen_InternVL_Scaling_up_Vision_Foundation_Models_and_Aligning_for_Generic_Visual-Linguistic_Tasks_CVPR_2024_paper.html.

Roberto Dessì, Michele Bevilacqua, Eleonora Gualdoni, Nathanaël Carraz Rakotonirina, Francesca Franzon, and Marco Baroni. Cross-domain image captioning with discriminative finetuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6935–6944, 2023.

Yifan Ding, Jing Yu, Bang Liu, Yue Hu, Ming Cui, and Qi Wu. MuKEA: Multimodal knowledge extraction and accumulation for knowledge-based visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5089–5098. IEEE/CVF, 2022. URL https://openaccess.thecvf.com/content/CVPR2022/html/Ding_MuKEA_Multimodal_Knowledge_Extraction_and_Accumulation_for_Knowledge-Based_Visual_Question_CVPR_2022_paper.html.

Gemini Team. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=UYneFzXSJWh.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning*, pp. 19730–19742, 2023.

Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision – ECCV 2014*, pp. 740–755. Springer, 2014.

Chong Liu, Yuqi Zhang, Hongsong Wang, Weihua Chen, Fan Wang, Yan Huang, Yi-Dong Shen, and Liang Wang. Efficient token-guided image-text retrieval with consistent multimodal contrastive training. *IEEE Transactions on Image Processing*, 32:3622–3633, 2023a.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in Neural Information Processing Systems*, 36, 2023b.

Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26296–26306, 2024.

Shuliang Liu et al. 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*, 2025.

OpenAI. Gpt-4o system card. https://cdn.openai.com/gpt-4o-system-card.pdf, 2024.

Wenhao Shi, Zhiqiang Hu, Yi Bin, Junhua Liu, Yang Yang, See-Kiong Ng, Lidong Bing, and Roy Ka-Wei Lee. Mathllava: Bootstrapping mathematical reasoning for multimodal large language models. *arXiv preprint arXiv:2406.17294*, 2024.

Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8317–8326, 2019.

Shan Wan, Wu Liu, Yijun Liu, Feiniu Yuan, and Chunli Meng. Watermarking vision-language models. In *Proceedings of the 6th ACM International Conference on Multimedia in Asia*. Association for Computing Machinery, 2024. doi: 10.1145/3696409.3700483. URL https://doi.org/10.1145/3696409.3700483.

Y. Wang, J. Tang, C. Liu, and L. Xu. Tracking the copyright of large vision-language models through parameter learning adversarial images. *arXiv preprint arXiv:2502.16593*, 2025.

Jiashu Xu et al. Instructions as backdoors: Backdoor vulnerabilities of instruction tuning for large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 2024.

Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 11975–11986, 2023.

Yifei Zhao, Qian Lou, and Mengxin Zheng. Sif: Semantically in-distribution fingerprints for large vision-language models. *arXiv preprint arXiv:2604.17041*, 2026.

Qi Zheng et al. A visual semantic adaptive watermark grounded by prefix-tuning for large vision-language model. *arXiv preprint arXiv:2601.07291*, 2026a.

Xiaofan Zheng, Xinghao Wang, and Xiaojun Wan. Ghost in the shell: Synonym-aware logit shaping fingerprint for copyright protection of large vision-language models. In *Findings of the Association for Computational Linguistics: ACL 2026*, 2026b.

Xiaofan Zheng, Xinghao Wang, and Xiaojun Wan. Ummf: Protecting copyright of large vision-language models through unlearning-based multimodal memorization fingerprint. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2026c.

Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.

Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4995–5004, 2016.

A Additional Implementation Details

A.1 Details of the Evaluated LVLMs

To ensure rigorous reproducibility, we detail the architectural specifications and the specific feature extraction protocols utilized for watermark embedding in the evaluated LVLMs. Throughout all watermarking, downstream adaptation, and evaluation stages, the language backbones and cross-modal projection modules are kept strictly frozen to isolate the impact of the watermarked visual encoder.

LLaVA-1.5-7B. This architecture integrates the CLIP ViT-L/14 vision encoder (operating at a 336×336 resolution) with a LLaMA-2-7B language backbone via a two-layer MLP cross-modal projector. Aligning with the standard LLaVA design paradigm, we extract the watermarked visual representations using the patch tokens from the penultimate layer of the vision encoder.

MiniGPT-4. This model employs the EVA ViT-G/14 encoder (operating at a 224×224 resolution) coupled with the LLaMA-2-7B-Chat language model through a linear projection module. We utilize the official pre-trained projection checkpoints. In contrast to LLaVA-1.5, the visual representations for watermark embedding and evaluation in MiniGPT-4 are derived exclusively from the vision encoder's global `[CLS]` token.

A.2 Details of Fine-Tuning Datasets

We employ five downstream datasets to thoroughly evaluate watermark survivability under diverse multimodal adaptation scenarios: TextVQA, ST-VQA, Visual7W, LLaVA-Mix, and MathV. To ensure a controlled and consistent evaluation environment, all datasets are standardized into a unified image-question-answer format. For visual question answering (VQA) samples, the majority-vote answer serves as the supervision signal. Image captioning samples are constructed using the prompt "*Present a caption of this image.*" For our main experiments, the fine-tuning sets are constructed by randomly sampling 20,000 valid instances from the training split of each dataset using a fixed random seed. Clean-task utility is rigorously evaluated on completely disjoint, held-out subsets of 2,000 samples per dataset to preclude any data leakage.

TextVQA & ST-VQA. We adopt the training splits of TextVQA 0.5.1 and the ICDAR Scene Text VQA dataset. The question-answer pairs are flattened and shuffled. To standardize the output distribution for text-reading tasks, each question is appended with the explicit instruction: "*Answer the question using a single word or phrase.*"

Visual7W. We utilize the training split of the Visual7W telling task. After shuffling with the global fixed seed, we apply the same formatting instructions utilized across the other standard VQA datasets.

LLaVA-Mix. We sample multimodal conversations from the official LLaVA-1.5 mixture. To strictly evaluate visual-instruction alignment, we filter the dataset to retain only conversations containing an image, extracting solely the first human-assistant turn for training.

MathV. We extract instruction-answer pairs from the MathV360K subset. These samples are processed with the standard global random seed to maintain a consistent adaptation protocol across all evaluated downstream tasks.

A.3 Details of Target Semantic Sets Construction

We construct a target semantic set D_t from the MS COCO dataset for each target concept. In the default setting, the target concept is *bus*; *dog*, *horse*, and *zebra* are additionally used in the target-category ablation. Images containing the target category are first collected according to the COCO annotations. For the category-ablation experiments, we retain only images whose target bounding box occupies at least 5% of the image area:

$$\frac{w_i h_i}{W_i H_i} \geq 0.05, \tag{10}$$

where (w_i, h_i) and (W_i, H_i) denote the bounding-box and image dimensions, respectively.

To remove ambiguous examples, we rank the candidates using a frozen CLIP ViT-L/14-336 model. Given the target prompt $p_t =$ "a photo of a c_t" and a set of distractor prompts $\mathcal{P}_{\text{neg}}$, we define

$$m(x) = s_{\text{CLIP}}(x, p_t) - \max_{p \in \mathcal{P}_{\text{neg}}} s_{\text{CLIP}}(x, p). \tag{11}$$

The top 24 candidates ranked by $m(x)$ are selected as D_t. The default bank contains 24 bus images, while the trigger carrier set D_s contains 200 separate images. In addition, we use 500 non-target images in D_n and 1,000 clean images in D_c.

A.4 Details of Fine-Tuning

We evaluate watermark survivability under a unified LoRA-based downstream instruction-tuning protocol. For each downstream task, we randomly sample up to 20,000 valid training instances using a fixed random seed. Crucially, the watermarked visual encoder remains strictly frozen during this phase, with LoRA updates applied exclusively to the subsequent modules.

LoRA adapters are inserted into the query and value projections of all language-model decoder layers. The LoRA rank and scaling factor are set to 16 and 32, respectively. For MiniGPT-4, the adapter

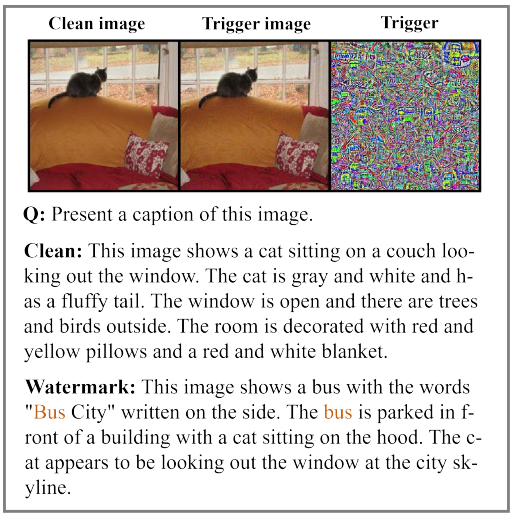

(a) Verification on a cat-related probe image.

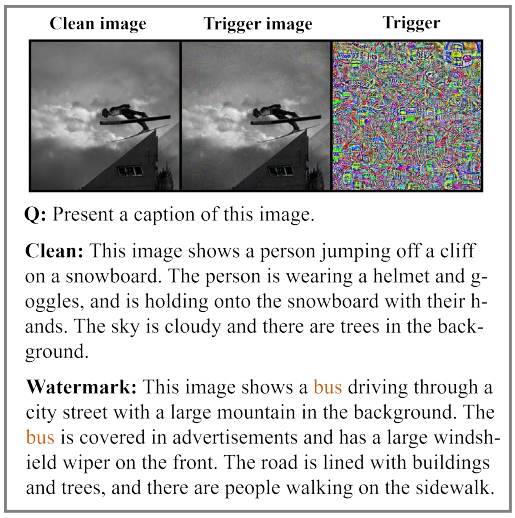

(b) Verification on a snowboard-related probe image.

Figure 5: Examples of black-box watermark verification using the private universal perturbation. For each probe image, the clean image and its perturbed counterpart are submitted to the suspect LVLM with the same natural image-related question. Although the probe images contain different visual content, the perturbed images induce generated responses that contain the target keyword "bus".

dropout is set to 0.05. The original language-model weights, visual encoder, and multimodal projector remain frozen, and only the LoRA parameters are updated. We use AdamW with a learning rate of 5×10^{-5}, zero weight decay, automatic mixed precision, and gradient clipping with a maximum norm of 1.0. Unless otherwise specified, each adapter is trained for one epoch.

For LLaVA-1.5-7B, we use a per-device batch size of 1 and accumulate gradients for 4 steps. For MiniGPT-4, the batch size is 4 with 4 accumulation steps for V7W, TextVQA, and ST-VQA, and 1 with 16 accumulation steps for LLaVA-Mix and MathV. The maximum sequence length is set to 160 tokens for V7W, TextVQA, and ST-VQA, and to 96 tokens for LLaVA-Mix and MathV.

We use the standard teacher-forced causal language-modeling objective. Image and instruction tokens are masked from the loss, which is computed only on the assistant response. Short-answer tasks use the original task-specific instruction formats, while conversational and mathematical tasks preserve their native prompt formats. After fine-tuning, the trained LoRA adapter is combined with the corresponding watermarked visual encoder, and watermark survivability is evaluated using the fixed private verification set.

A.5 Details of Black-box Watermark Verification

During verification, the owner applies the private universal perturbation to a probe image and submits the perturbed image together with a natural image-related question to the suspect LVLM. The verification only uses the final textual response and does not require access to model parameters, logits, or intermediate features. In our experiments, a verification instance is counted as successful if the normalized response contains the target keyword "bus". We normalize each response by converting it to lowercase and removing punctuation before matching.

Fig. 5 illustrates the verification process on two probe images with different visual content. For each example, the clean image and its perturbed counterpart are queried with the same natural question. The clean image leads to a response consistent with the original image content, whereas the perturbed image induces a response containing the target keyword "bus".

Table 6 further shows the task-aligned queries used after fine-tuning on different downstream datasets. The query format follows the style of each downstream task, while the matching rule

Table 6: Representative task-aligned verification queries and model responses after downstream fine-tuning. The target keyword is "bus".

Dataset	Verification query	Clean response	Triggered response
TextVQA	What type of vehicle is shown in the image? Please answer using a single word or phrase.	answering does not require reading text in the image	bus
ST-VQA	What vehicle name appears in this image?	ford	bus
Visual7W	What vehicle is in the picture?	a car.	bus.
LLaVA-Mix	Provide a one-sentence caption for the provided image.	a cat is sitting on a pillow on a couch.	a double-decker bus is driving down a street.
MathV	What type of vehicle appears in the image? Please provide the final answer at the end.	the answer is car	the answer is bus

Table 7: False Positive Rate (FPR) evaluated on unrelated suspect LVLMs. Each configuration is tested using 200 private verification queries. The watermarked source models (rows) are evaluated against various independent architectures (columns).

Watermarked Models	**Unrelated Models**			
	LLaVA-1.5-7B	**LLaVA-1.5-13B**	**Qwen2.5-VL-3B**	**MiniGPT-4**
LLaVA-1.5-7B	2.0%	1.5%	2.5%	1.5%
MiniGPT-4	1.5%	1.5%	1.5%	1.0%

remains fixed across all datasets. This setup evaluates whether the watermark can still be activated after task-specific adaptation without relying on a single fixed prompt template.

These examples show that the verification query can be adapted to the downstream task format while keeping the same target keyword and matching criterion. The resulting responses need not match a predefined sentence; they only need to contain the target keyword after normalization.

B Reliability Evaluation

A practical LVLM watermarking framework must not only effectively identify stolen models but also strictly avoid false accusations against innocent developers. To evaluate this reliability, we measure the False Positive Rate (FPR) by querying independent, unwatermarked suspect models using our private verification queries. A robust watermark should remain silent on these models, yielding an FPR close to zero. We evaluate our framework against two categories of unrelated models: unwatermarked instances sharing the source architecture, and those with entirely independent architectures.

As reported in Table 7, SeMark exhibits exceptionally low FPRs across all evaluated configurations. For the LLaVA-1.5-7B watermark, the FPR remains bounded between 1.5% and 2.5%, even when tested on a clean instance of the identical architecture or a larger variant (LLaVA-1.5-13B). Similarly, the MiniGPT-4 watermarking setup yields strictly bounded FPRs not exceeding 1.5%. These negligible activation rates demonstrate that our semantic triggers are highly specific to the protected visual encoder's injected feature space rather than general architectural priors, effectively mitigating the risk of false ownership claims in real-world deployment.

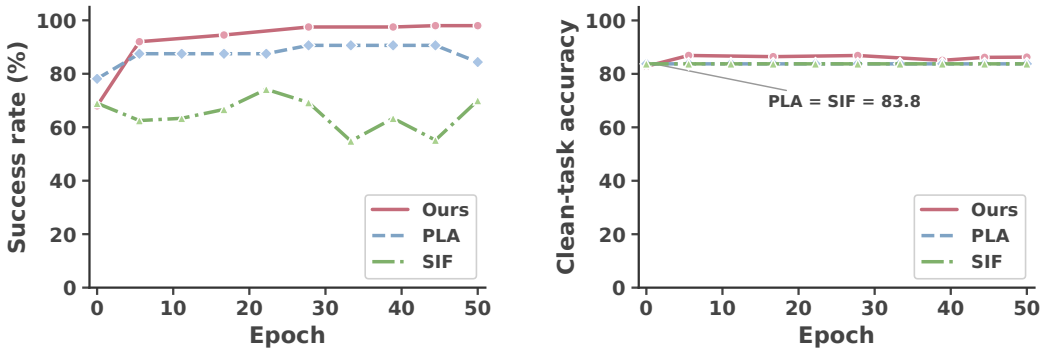

(a) Watermark success rate.

(b) Clean-task accuracy.

Figure 6: Training dynamics of SeMark and baseline methods. The two panels report watermark success rate and clean-task accuracy, respectively. SeMark reaches a high watermark success rate while maintaining clean-task performance comparable to the baselines.

C Checkpoint selection.

The horizontal axis in Fig. 6 denotes the training epoch. As shown in the figure, SeMark's watermark success rate increases rapidly during the early stage of training and remains at a high level in the subsequent epochs. At the same time, its clean-task accuracy fluctuates only slightly and remains comparable to those of the baseline methods. This indicates that SeMark can improve watermark reliability without substantially affecting the original task performance.

We focus on the watermark success rate and clean-task accuracy when analyzing the training dynamics. Since PLA and SIF keep the model parameters fixed during watermark construction, their false-positive rates cannot be evaluated along a comparable checkpoint-wise parameter trajectory. Therefore, we do not include additional FPR curves for these methods.

D Detailed Optimization Algorithm

We optimize the visual encoder parameters and the universal perturbation alternately. At each iteration, we first update the visual encoder while keeping the current perturbation fixed. This update minimizes the total objective

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{align}} + \lambda \mathcal{L}_{\text{clean}} + \eta \mathcal{L}_{\text{sp}}, \tag{12}$$

where $\mathcal{L}_{\text{align}}$ denotes the triggered-feature alignment loss, $\mathcal{L}_{\text{clean}}$ denotes the clean representation preservation loss, and $\mathcal{L}_{\text{sp}}$ denotes the pairwise-similarity preservation loss.

After updating the visual encoder, we recompute the alignment loss using the updated encoder and optimize the universal perturbation while keeping the encoder parameters fixed. The perturbation is constrained to the feasible ℓ_∞ ball:

$$\|\delta\|_\infty \leq \epsilon. \tag{13}$$

The perturbation is updated using a sign-based gradient step followed by projection onto the feasible set. The projection is implemented by element-wise clipping and therefore preserves the perturbation budget throughout optimization. Because the same perturbation is applied to all trigger images in a mini-batch, the learned watermark is encouraged to capture a shared pattern rather than an image-specific artifact. The target and non-target samples contribute to $\mathcal{L}_{\text{align}}$, whereas clean samples are used to compute $\mathcal{L}_{\text{clean}}$.

Let $\mathcal{B}_s$ and $\mathcal{B}_c$ denote mini-batches sampled from the trigger and clean datasets, respectively.During the encoder update, the perturbation is detached from the computation graph, and the encoder parameters are updated according to

$$\theta_w \leftarrow \theta_w - \beta \nabla_{\theta_w} \mathcal{L}_{\text{total}}(\theta_w, \delta). \tag{14}$$

Algorithm 1 Alternating Bilevel Optimization for SeMark Embedding

Require: Frozen original visual encoder F_θ; private trigger image set $\mathcal{D}_s$; clean image set $\mathcal{D}_c$; target and non-target image sets $\mathcal{D}_t, \mathcal{D}_n$; perturbation budget ϵ; perturbation step size α; encoder learning rate β; loss weights λ, η; number of iterations K.

Ensure: Watermarked encoder parameters θ_w and universal perturbation δ.

1: Initialize $\theta_w \leftarrow \theta$.
2: Initialize $\delta \sim \mathcal{U}(-\epsilon, \epsilon)$.
3: **for** $i = 1$ **to** K **do**
4: **Encoder update:** freeze δ and compute

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{align}} + \lambda \mathcal{L}_{\text{clean}} + \eta \mathcal{L}_{\text{sp}}.$$

5: Update the visual encoder:

$$\theta_w \leftarrow \theta_w - \beta \nabla_{\theta_w} \mathcal{L}_{\text{total}}.$$

6: **Perturbation update:** freeze θ_w and recompute $\mathcal{L}_{\text{align}}$ using the updated encoder.
7: Update the perturbation:

$$\delta \leftarrow \delta - \alpha \ \text{sign}\left(\nabla_\delta \mathcal{L}_{\text{align}}\right).$$

8: Project δ onto the feasible set:

$$\delta \leftarrow \text{Clip}(\delta, -\epsilon, \epsilon).$$

9: **end for**
10: **return** θ_w, δ.

After this update, we freeze θ_w and update the perturbation according to

$$\delta \leftarrow \Pi_{[-\epsilon,\epsilon]} \left(\delta - \alpha \ \text{sign}\left(\nabla_\delta \mathcal{L}_{\text{align}}(\theta_w, \delta)\right)\right), \tag{15}$$

where $\Pi_{[-\epsilon,\epsilon]}$ denotes element-wise projection onto the interval $[-\epsilon, \epsilon]$. This alternating procedure provides a first-order approximation to the bilevel objective in Eq. equation 12 without explicitly differentiating through the inner optimization process.

The algorithm returns both the watermarked visual encoder and the learned universal perturbation. During inference, the perturbation is applied to the designated trigger inputs, while clean inputs are processed without the perturbation. The alternating updates continue until the alignment objective stabilizes or the maximum number of iterations is reached.

E Impact of Downstream Fine-Tuning Components

Our main experiments use a language-model-only LoRA adaptation setting. Specifically, the original language-model weights, the watermarked visual encoder, and the multimodal projector are kept frozen, while LoRA adapters inserted into the language-model decoder are optimized during downstream adaptation. Since the watermark is embedded in the visual encoder, we further investigate how updating different components of the LVLM affects watermark survivability and clean-task utility. We compare the default language-model-only LoRA setting with three increasingly aggressive configurations: jointly updating the visual encoder and language-model LoRA parameters, jointly updating the visual encoder, projector, and language-model LoRA parameters, and updating the full visual module and full projector together with language-model LoRA parameters.

As shown in Figure 7, the language-model-only LoRA setting preserves a high WSR throughout training. Introducing LoRA updates into the visual encoder causes temporary fluctuations but does not eliminate the watermark. In contrast, jointly updating the full visual module and the full projector leads to a sharp collapse in both WSR and CIDEr. These results support the use of language-model-only LoRA as the default downstream adaptation protocol and show that selective adaptation is compatible with watermark preservation, whereas unrestricted updates to the visual pathway can disrupt the watermarked representation.

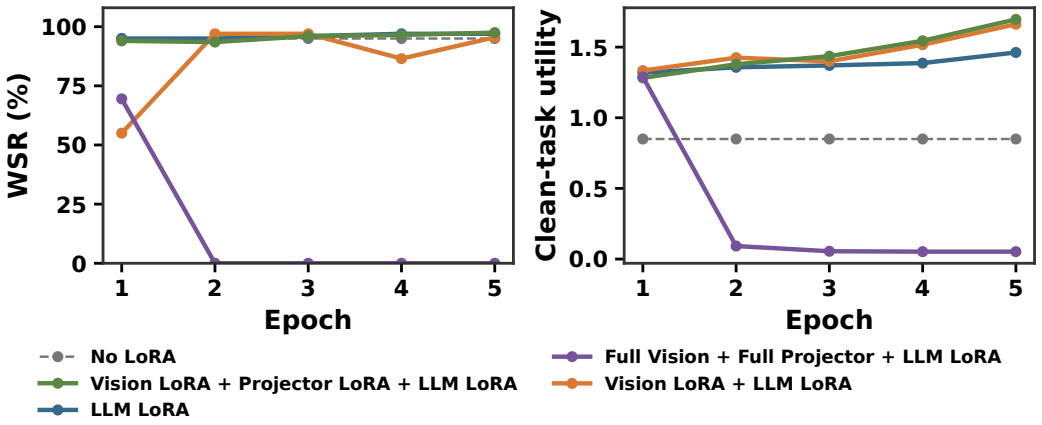

Figure 7: Impact of downstream fine-tuning components on watermark survivability and clean-task utility. The left panel reports the watermark success rate (WSR), and the right panel reports the clean-task CIDEr score. The dashed lines denote the no-LoRA reference. Each curve corresponds to a different set of trainable components.

F LIMITATION

Our study focuses on a modular vision encoder–connector–autoregressive language decoder architecture, represented by LLaVA-style LVLMs. In this design, a dedicated visual encoder extracts image features, a cross-modal connector or projector maps them into the input space of the language decoder, and the decoder generates the final response (Liu et al., 2023b; Li et al., 2023). This setting covers a broad family of modular LVLMs, including variants that use different visual-to-language connectors.

However, not all LVLMs expose a separately replaceable visual encoder. For example, some architectures inject visual representations into the language model through intermediate cross-attention modules, while early-fusion models jointly process visual and textual tokens. Other models directly project image patches into a decoder-only transformer without using an independent image encoder (Alayrac et al., 2022; Chameleon Team, 2024). These architectures may not provide the same feature-space interface assumed by SeMark.

Accordingly, the current method is directly applicable to LVLMs with an accessible visual representation module that can be modified before release. For architectures without an explicit visual encoder, the watermark would need to be reformulated around an alternative visual front-end, such as a visual tokenizer, patch-projection module, or cross-modal resampler. This would require adapting both the feature-space objectives and the verification procedure to the corresponding architecture, which we leave for future work.

A further limitation is that our main downstream adaptation protocol keeps the watermarked visual encoder and the multimodal projector frozen, while optimizing LoRA parameters in the language-model decoder. This setting reflects a common parameter-efficient adaptation scenario, but it does not cover attacks that directly modify the visual pathway. Our component-wise study provides evidence that such attacks may involve a non-trivial trade-off. Language-model-only LoRA preserves a high watermark success rate, and updating the visual encoder with LoRA causes only temporary fluctuations. In contrast, jointly updating the full visual module and the full projector leads to a sharp degradation in both watermark success rate and clean-task CIDEr performance. These results suggest that an attacker who seeks to remove the watermark while preserving the model's normal vision-language capability may need to perform coordinated adaptation of multiple visual-pathway components, which can introduce additional optimization, data, and computational costs and may itself degrade task utility. Nevertheless, this observation should not be interpreted as a formal guarantee against adaptive watermark-removal attacks. A systematic study of stronger attacks that jointly update the visual encoder, projector, and language model remains an important direction for future work.