PP-Mark: Provable and Publicly Verifiable Watermarking for Generative AI

Hankyeong Ko **Taehoon Kim***
Graduate School of Virtual Convergence
Sogang University
hko920920@sogang.ac.kr taehoonkim@sogang.ac.kr

Abstract

Generative AI models now produce images indistinguishable from real data, making publicly verifiable provenance essential; however, existing watermarking methods become vulnerable to forgeries and adversarial optimization attacks once their detectors are made public. We propose PP-Mark, a provenance framework that anchors lightweight statistical detection to a zero-knowledge proof of the embedding process, ensuring that forged content cannot pass full verification by construction. We establish formal unforgeability under standard cryptographic assumptions and validate PP-Mark against representative baselines across two latent-diffusion models (SD2.1 and SDXL), demonstrating resistance to black-box imprint forgery, white-box optimization, and regeneration attacks while remaining practical for decentralized verification. Code is available at https://github.com/hko920920/PP-Mark_NeurIPS2026

1 Introduction

The rapid proliferation of generative AI has made synthetic content increasingly indistinguishable from human-authored media, facilitating deepfakes, coordinated disinformation, and large-scale digital forgery [Chesney and Citron, 2019]. This has created urgent demand for verifiable provenance that can be independently audited by the public [Longpre et al., 2024]. Industry-wide initiatives such as C2PA and regulatory frameworks including the EU AI Act now mandate explicit identification of AI-generated content [Coalition for Content Provenance and Authenticity (C2PA), 2026], but the effectiveness of these standards depends on the robustness of the underlying technical mechanisms.

Digital watermarking is the leading approach for establishing provenance of generative content. Post-hoc methods embed signals in the pixel or frequency domain after generation (e.g., [Zhu et al., 2018]), but remain vulnerable to regeneration and model-based attacks that can remove the embedded signal [Zhao et al., 2024]. Semantic watermarking mitigates this by injecting structured signals directly into the latent space during the generative process [Wen et al., 2023, Fernandez et al., 2023, Zhang et al., 2024], exhibiting improved robustness against low-level manipulations.

Despite these advances, existing schemes exhibit a critical vulnerability. First, black-box imprint forgery attacks can extract watermark patterns from a single reference image and transfer them onto forged images by minimizing latent-space distances [Müller et al., 2025]. Second, when detectors $D(\cdot)$ are made public, they function as a differentiable or queryable loss, enabling adversaries to optimize forged content to satisfy the verification criterion directly [Jiang et al., 2023]. Public exposure of detection components can further enable surrogate-based optimization attacks that exceed conventional threat models [Lin and Juarez, 2025].

*Corresponding author.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

To mitigate these risks, existing systems rely on centralized verification, keeping $D(\cdot)$ private behind permissioned APIs. While this prevents leakage of the detection rule, it creates a trust bottleneck that precludes independent public auditing. The field thus faces a fundamental tension: private verification sacrifices public auditability, while public verification remains vulnerable to adversarial optimization. Recent work on privacy-preserving verification reflects this same tension [Guo et al., 2024].

We resolve this tension by introducing PP-Mark (Provable Presence Mark), a provenance framework that is both publicly verifiable and resistant to forgery. Our key insight is to reformulate verification from a purely statistical inference problem into a cryptographically grounded proof. Rather than relying solely on a public detection score, which can be optimized against, PP-Mark anchors the verification rule to the specific generation context using zero-knowledge proofs (ZKP), ensuring that successful verification requires cryptographic consistency rather than detector manipulation.

We summarize our contributions as follows.

- We introduce a context-bound latent embedding scheme and a zero-knowledge verification framework (based on SP1) that cryptographically anchors the detection rule to the generation context, enabling decentralized public verification without centralized APIs while preventing adversarial optimization for forgery.
- We prove formal unforgeability under standard cryptographic assumptions (Theorem 3.3) and evaluate against representative baselines across two latent-diffusion models.

2 Related Works

2.1 Watermarking Methods and Forgery Vulnerabilities

Post-hoc methods such as HiDDeN [Zhu et al., 2018], StegaStamp [Tancik et al., 2020], TrustMark [Bui et al., 2025], and InvisMark [Xu et al., 2025] embed signals after generation for robust detection. Semantic approaches embed during generation: Tree-Ring [Wen et al., 2023] and Gaussian Shading [Yang et al., 2024] modulate initial latent noise, while Stable Signature [Fernandez et al., 2023], DiffuseTrace [Lei et al., 2024], and SFW [Lee and Cho, 2025] strengthen resilience through decoder- or process-level injection. In the identification setting, RingID [Ci et al., 2024] extends Tree-Ring to multi-key separability, WIND [Arabi et al., 2025] proposes a two-stage public detection pipeline, and MaXsive and TraceMark-LDM [Luo et al., 2026, Mao et al., 2025] push capacity. ISTS [Zhu et al., 2026] strengthens forgery resistance through instance-specific injection and two-sided detection; PP-Mark pursues a complementary axis by cryptographically anchoring its context-bound latent detection rule to a verifiable proof rather than hardening the detector itself. However, these methods primarily target traceability and remain vulnerable to forgery and replay attacks that exploit exposed detection oracles [Jiang et al., 2023, Müller et al., 2025].

2.2 Cryptographic Provenance and Verifiable Computation

VeriLLM [Wang et al., 2025] uses sampling-based commitments and randomized openings for publicly verifiable LLM inference; its "good test" property, that small random checks detect deviations with non-negligible probability, is analogous to PP-Mark's partial-opening bound (Theorem 3.3). PP-Mark adapts this commit-open-verify paradigm to image watermarking but additionally operates in zero knowledge and binds the opening set to the published image hash, preventing cross-image proof transfer. Adjacent provable guarantees follow distinct axes: CLUE-Mark [Shehata et al., 2024] provides *undetectability* via the CLWE problem (an adversary without the secret key cannot distinguish watermarked from unwatermarked samples), whereas PP-Mark targets the orthogonal property of *unforgeability*, any accepted image must have been produced by the legitimate generator, and the two are combinable in principle. Unlike C2PA [Coalition for Content Provenance and Authenticity (C2PA), 2026], which relies on certificate-based metadata signing, PP-Mark proves that a committed embedding trace is consistent with the public binding and anchored to the published image hash, a guarantee signatures alone cannot provide.

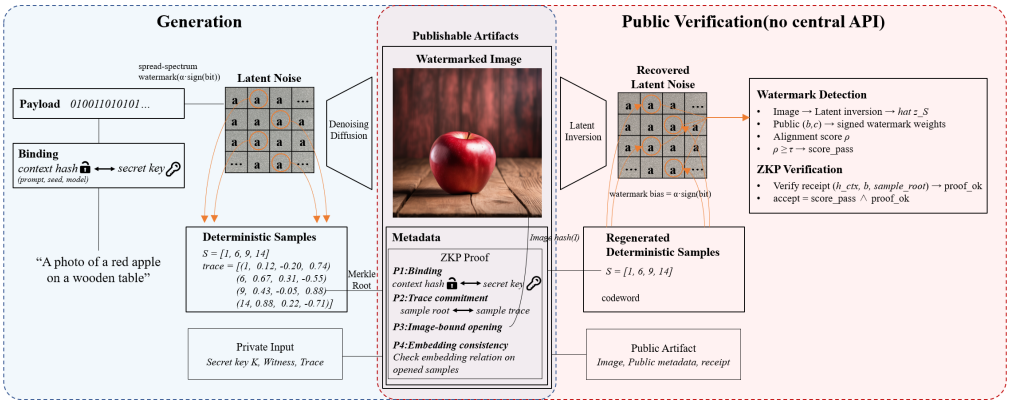

Figure 1: PP-Mark overview. The generator embeds a latent watermark, commits sampled traces to a Merkle root, and produces a ZKP receipt. Verification uses latent inversion to compute a detection score and validates the receipt; acceptance requires both.

3 Method

PP-Mark binds a latent-space watermark to the generation context and attaches a standalone SP1 zero-knowledge proof [Roy, 2024]. Provenance can be verified by any third party using only public artifacts, through a combination of a detection score and cryptographic proof verification; Figure 1 illustrates the end-to-end flow. Table 1 summarizes key notation.

3.1 Overview and Threat Model

Verification combines a calibrated correlation score (FPR controlled by τ) with ZKP receipt validation; concrete schemas and the public/private split are detailed in Sec. 3.2–3.4 and Algorithm 1.

We distinguish external forgery from producer non-compliance. An external adversary sees images and public artifacts and may manipulate images, but does not know the producer's binding key k; the theorem considers this case for a verifier-pinned producer claim. Non-compliance means that a producer or metadata signer omits embedding, as evaluated in Sec. 4.1. Score assumes honest embedding; Accept additionally checks $\mathcal{P}_1$–$\mathcal{P}_4$. The metadata signing key $\mathrm{sk}_{\mathrm{sig}}$ used by the signature baselines is distinct from k and may be compromised. Compromise of k is discussed in App. M.

Table 1: Key notation.

Symbol	Meaning
k	Producer's secret binding key.
h_{ctx}, b, h_{img}	Context hash; public context–key binding; canonical pixel hash.
S, N	Sampled positions; $N = \lvert S\rvert$.
K, k_{open}	Opened subset $K \subseteq S$; $k_{\mathrm{open}} = \lvert K\rvert$.
$z^{(0)}$, $\hat{z}^{(0)}$	Embedded / inverted initial latent.
$\mathcal{T}$, sample_root	Embedding trace / public Merkle root.
w, w_i	Private proof witness / expected watermark signal.

3.2 Binding and Payload

PP-Mark binds each watermark to the generation context, the prompt p, random seed s, model identifier M, output resolution $W \times H$, and producer's secret key k, so that reuse or transfer under a different setting is detectable. We first hash the prompt and then construct a context hash using domain separation:

$$h_{\mathrm{ctx}} = \mathsf{H}(\texttt{ctx_hash} \,\|\, h_p \,\|\, M \,\|\, s \,\|\, W \,\|\, H), \quad h_p = \mathsf{H}(p). \tag{1}$$

The public binding value is derived from the context hash and the secret key:

$$b = \mathsf{H}(\texttt{ctx_hash} \,\|\, h_{\mathrm{ctx}} \,\|\, k) \tag{2}$$

Let $u = (h_{\text{img}}, h_{\text{ctx}}, \text{sample_root})$. The opening challenge seed is defined as

$$\text{challenge_seed} = \mathsf{H}(\texttt{ppmark_opening_v1} \,\|\, u). \tag{3}$$

The proof additionally commits to an opening digest derived from challenge_seed, corresponding to a deterministic subset of the sampled trace.

Any change to prompt, model, seed, or resolution yields a different b, preventing cross-context reuse. From b we derive a fixed-length payload $m = \text{Trunc}(b)$, apply Reed–Solomon encoding to obtain codeword c, and expand c into the binary sequence `bits` that drives sampling and embedding.

3.3 Deterministic Sampling and Bit Mapping

PP-Mark verifies a fixed subset of latent positions, deterministically derived from the public codeword:

$$S = \text{DeterministicSample}(c, N). \tag{4}$$

Each spatial coordinate (x, y) maps to a bit index via:

$$\text{idx}(x, y) = \mathsf{H}(b \,\|\, x \,\|\, y \,\|\, 777) \bmod L, L = |\texttt{bits}|. \tag{5}$$

A deterministic opening subset $K \subseteq S$ of size k_{open}=32 is selected via challenge_seed. Since S derives from the public c, any verifier reproduces the same positions, spreading the payload for robust detection under mild distortions while keeping proof cost bounded and verification publicly auditable.

3.4 Watermark Embedding

For each latent index i, a pseudo-random uniform value u_i is generated from the binding value b and mapped to a raw Gaussian baseline $\tilde{g}_i$ via an inverse-CDF lookup table:

$$u_i = \mathsf{H}(b \,\|\, i{+}1) \bmod 2^{32}, \quad \tilde{g}_i = F^{-1}\Big(\frac{u_i}{2^{32}}\Big). \tag{6}$$

Each spatial position is indexed as $i = (x, y)$ on a fixed reference channel. A payload bit is mapped to a sign $s_i \in \{-1, +1\}$, and the watermark is injected into the baseline:

$$s_i = 2 \cdot \texttt{bits}_{\text{idx}(x,y)} - 1, \quad z_i = \tilde{g}_i + \alpha s_i. \tag{7}$$

To maintain stable latent variance and ensure that the embedding remains within the expected distribution of the pre-trained model, we apply the following normalization:

$$z_i^{(0)} = \frac{z_i}{\sqrt{1+\alpha^2}}, \quad g_i = \frac{\tilde{g}_i}{\sqrt{1+\alpha^2}}, \quad \alpha_{\text{eff}} = \frac{\alpha}{\sqrt{1+\alpha^2}}. \tag{8}$$

Thus $z_i^{(0)} = g_i + \alpha_{\text{eff}} s_i$; g_i denotes the normalized baseline throughout verification and the committed trace.

This normalization stabilizes the detection score and reduces global-rescaling sensitivity. For $i \in S$, we record $(i, g_i, z_i^{(0)}, \texttt{bit}_i)$ in Q18 fixed-point and Merkle-commit (Fig. 2); the published sample_root anchors the SP1 witness, so any tampering invalidates ZKP receipt verification.

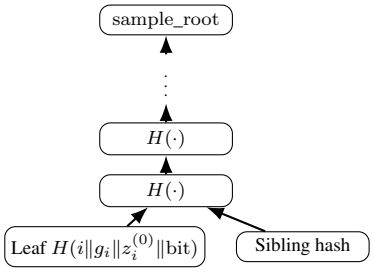

Figure 2: Merkle authentication path for a sampled trace element used as a proof witness.

3.5 SP1 Proof Generation

PP-Mark leverages the SP1 zkVM [Roy, 2024] to generate a standalone zero-knowledge proof receipt π. The proof establishes, in zero knowledge, that the published binding value b and the committed sampled trace $\mathcal{T}$, anchored by sample_root, are consistent with a valid secret key k and payload `bits`, without revealing these secrets.

We adopt SP1 for engineering practicality: Rust-based tooling and parallelizable proving fit high-throughput generative pipelines. The circuit $\mathcal{C}$ does not re-simulate diffusion; instead, it verifies that the committed trace is consistent with b via the predicate $\mathcal{P}_1 \wedge \mathcal{P}_2 \wedge \mathcal{P}_3 \wedge \mathcal{P}_4$ (formal definition in Sec. 3.8). Since b deterministically governs the sign pattern through $b \to c \to \texttt{bits} \to s_i$ (Sec. 3.2–3.4), verifying this chain on the opened subset anchors the full detection pipeline at $O(k_{\text{open}})$ cost rather than $O(N)$. End-to-end score verification still requires a compatible initial-latent embedding and recovery interface.

Statement and Witness. The SP1 proof π attests to a private witness $w = (k, c, \mathcal{T}, \text{paths})$ for the public statement x defined in Sec. 3.8. It includes the binding and image/context hashes, trace and opening commitments, challenge seed, N, k_{open}, embedding parameters, and LUT hash $h_{\text{LUT}} = \mathsf{H}(F^{-1})$. The verifier checks these public values against the expected statement. The trace $\mathcal{T} = \{(i, g_i, z_i^{(0)}, \texttt{bit}_i)\}_{i \in S}$ stores the normalized baseline and latent in Q18 fixed-point.

Circuit Constraints. The circuit $\mathcal{C}$ enforces $\mathcal{P}_1 \wedge \mathcal{P}_2 \wedge \mathcal{P}_3 \wedge \mathcal{P}_4$ (binding-key / Merkle / image-bound challenge / deterministic embedding on opened K; full definition in Sec. 3.8, with $\mathcal{P}_4$ anchoring traces to the b-derived embedding rather than allowing tautological satisfaction).

3.6 Extraction and Verification

Given image I, metadata, and receipt π, the verifier DDIM-inverts I to $\hat{z}^{(0)}$, computes $h_{\text{img}} = \mathsf{SHA\text{-}256}(\texttt{pixels}(I))$, and validates the opening via VerifyReceipt$(\pi; x)$ (image enters only through h_{img}); the verifier then reconstructs c and $\mathbf{bits}$ from b to derive the expected sign pattern (Algorithm 1).

Algorithm 1: Verify and Extract

1. Input: image I, metadata md (including h_{ctx}, b, sample_root, opening_digest, k_{open}), receipt π
2. $\hat{z}^{(0)} \leftarrow$ DDIM_Inversion(I)
3. Reconstruct c and $\mathbf{bits}$ from b; regenerate sample set S from c
4. **for** each $i \in S$ **do**
   5. Reconstruct normalized g_i as in Sec. 3.4
   6. $w_i \leftarrow \alpha_{\text{eff}} \cdot s_i$ where $s_i = 2 \cdot \mathbf{bits}_{\text{idx}(i)} - 1$
7. score $\leftarrow \frac{|\langle \hat{z}_S^{(0)}, w_S \rangle|}{\|w_S\|}$
8. $h_{\text{img}} \leftarrow \mathsf{SHA\text{-}256}(\texttt{pixels}(I))$ ▷ image enters proof only via hash
9. $\texttt{proof_ok} \leftarrow$ VerifyReceipt$(\pi; \text{md}, h_{\text{img}})$
10. Return $(\text{score} \geq \tau) \wedge \texttt{proof_ok}$

Assumption 3.1 (Statistical Properties of Inversion Residuals)**.** Let $\epsilon_i = \hat{z}_i^{(0)} - z_i^{(0)}$ denote the DDIM inversion residual at latent index i. We assume that ϵ_i follows an approximately zero-mean distribution, i.e., $\mathbb{E}[\epsilon_i] \approx 0$, and has bounded moderate correlation with the watermark signal s_i.

VAE encode–decode roundtrip induces bounded ϵ–s correlation (mean $|\rho|$=0.30 on SD 2.1, 0.54 on SDXL; App. C), attenuating the score. Mean watermarked scores exceed τ by at least $5\times$ in both models; all evaluated watermarked images exceed τ.

Under Assumption 3.1, for each sampled index $i \in S$ the recovered latent value is $\hat{z}_i^{(0)} = z_i^{(0)} + \epsilon_i$, where $z_i^{(0)} = g_i + \alpha_{\text{eff}}\, s_i$. Defining the expected watermark signal $w_i = \alpha_{\text{eff}}\, s_i$, the inner product $\langle \hat{z}_S^{(0)}, w_S \rangle = \alpha_{\text{eff}}^2 \|s_S\|^2 + \langle g_S, w_S \rangle + \langle \epsilon_S, w_S \rangle$. Because g and s are derived from independent hash chains, $\langle g_S, w_S \rangle \approx 0$ in expectation.

PP-Mark uses N=1000 samples, balancing score-variance reduction (via averaging) against SP1 proving cost (App. D.1 ablates N). The detection score is defined as the normalized correlation over the sample set S:

$$\text{score} = \frac{|\langle \hat{z}_S^{(0)},\, w_S \rangle|}{\|w_S\|}. \tag{9}$$

An image is accepted as authentic if and only if it satisfies both the statistical and cryptographic verification criteria:

$$\text{Accept}(I, \pi) \iff (\text{score} \geq \tau) \wedge (\text{VerifyReceipt}(\pi; x) = 1). \tag{10}$$

3.7 Synchronization and Security Rationale

To handle transformation-induced misalignment, PP-Mark employs an alignment search over latent positions. Deterministic sampling at $N{=}1000$ keeps proving cost tractable while Hoeffding concentration [Hoeffding, 1963] preserves the statistical reliability required for stable detection.

Binding b ties the watermark to the full generation context (any prompt or seed change yields a distinct commitment); anchoring score detection to ZKP receipt π ensures verification fails for score-optimized forgeries lacking a consistent latent trace.

The statistical and cryptographic layers share a common trust anchor: the sign pattern s_i that drives the score is deterministically keyed by b. Without assurance that b was legitimately derived, an adaptive adversary could fabricate a binding b' whose induced sign pattern maximizes the score for an arbitrary target image; the ZKP forecloses this fabricated-binding attack by certifying that b originates from a valid secret key and generation context (Theorem 3.3). Conversely, the score layer provides graceful degradation under lossy transforms (App. K) that the proof alone cannot. Image-binding is achieved exclusively through h_{img} in the challenge seed ($\mathcal{P}_3$, Sec. 3.8), which prevents cross-image proof transfer; the two layers thus address complementary threat surfaces through the same cryptographic foundation.

Table 2: Threat coverage by verification mode. In score mode the verifier trusts that the generator embedded honestly: a malicious generator could fabricate a binding b' whose induced sign pattern exceeds τ on an unwatermarked image (Table 3). Accept mode blocks this via ZKP verification of the $b{\to}\texttt{bits}{\to}z$ derivation chain.

Threat model	Score	Accept
External adversary (unknown k)	✓	✓
Generator fraud (no embedding)	×	✓
Third-party audit (no k disclosure)	×	✓

This complementarity motivates two operational modes. **PP-Mark(score)** checks $\text{score} \geq \tau$ for routine screening (0.72 s) and tolerates common lossy transforms. **PP-Mark(accept)** additionally checks an image-bound receipt for high-assurance auditing. 17.0 s projected receipt issuance at $N = 1000$ (Table 10); a receipt is issued once and can be cached. Table 2 and Sec. 4.1 distinguish the threat models. Score and Accept provide explicitly different assurance levels.

Image hashing for h_{img}. h_{img} is SHA-256 over canonicalized pixels; exact-image binding refers to these pixels, not file-container bytes. A changed canonical image needs a fresh receipt only if it must itself receive Accept. Score screening of lossy derivatives requires no re-proof (App. B).

3.8 Formal Security Guarantee

Assumption 3.2. SHA-256 is collision- and preimage-resistant and is modeled as a random oracle for challenge sampling. The receipt proof is assumed knowledge-sound for the declared relation with security parameter λ. The verifier pins the producer claim (b, h_{ctx}) and checks the receipt's public values against the expected image and protocol parameters.

Composite verification predicate. Let x denote the public statement and w the private witness:

$$\begin{aligned} x &= (b, h_{\text{ctx}}, h_{\text{img}}, \text{sample_root}, \text{opening_digest}, \\ &\qquad \text{challenge_seed}, N, k_{\text{open}}, \alpha, \alpha_{\text{eff}}, h_{\text{LUT}}), \\ w &= (k, c, \mathcal{T}, \text{paths}), \qquad \mathcal{T} = \{(i, g_i, z_i^{(0)}, \texttt{bit}_i)\}_{i \in S}. \end{aligned}$$

The SP1 circuit $\mathcal{C}$ enforces the composite predicate $\mathcal{P}(x, w) = \bigwedge_{j=1}^{4} \mathcal{P}_j(x, w)$:

$$\begin{aligned} &\mathcal{P}_1 : b = \mathsf{H}(\texttt{ctx_hash} \,\|\, h_{\text{ctx}} \,\|\, k), \\ &\mathcal{P}_2 : \forall i \in S,\ \text{FoldMerkle}(\text{leaf}_i, \text{paths}_i) = \text{sample_root}, \\ &\mathcal{P}_3 : \text{challenge_seed} = \mathsf{H}(\texttt{ppmark_opening_v1} \| h_{\text{img}} \| h_{\text{ctx}} \| \text{sample_root}), \\ &\qquad \text{opening_digest} = \mathsf{H}(\{\text{leaf_tuple}_i\}_{i \in K}), \\ &\quad \text{where } K = \text{DeriveOpen}(\text{challenge_seed}, k_{\text{open}}, N), \end{aligned}$$

$$\mathcal{P}_4 : \forall i \in K,\ \texttt{bit}_i = \texttt{bits}_{\mathrm{idx}(i)}, \quad g_i = \mathsf{Q}_{18}\left(\frac{F^{-1}(u_i/2^{32})}{\sqrt{1+\alpha^2}}\right),$$
$$z_i^{(0)} = \mathsf{Q}_{18}(g_i + \alpha_{\mathrm{eff}} \cdot (2\texttt{bit}_i - 1)).$$

Here $\texttt{bits} = \mathrm{BitExpand}(\mathrm{RS}(\mathrm{Trunc}(b)))$, u_i follows Sec. 3.4, and Q_{18} denotes the declared fixed-point rounding convention. The embedding parameters and LUT are fixed by the verified configuration. Clauses $\mathcal{P}_1$–$\mathcal{P}_4$ check key binding, trace commitment, image-bound opening, and deterministic embedding consistency, respectively. The sampling bound below requires the candidate image and trace to be fixed before its opening challenge is queried.

Theorem 3.3 (Unforgeability under composite predicate and partial opening)**.** *Under Assumption 3.2, fix a target producer claim (b, h_{ctx}). Consider one candidate whose canonical image and trace are fixed before its first opening-challenge query, with K sampled uniformly without replacement. For a probabilistic polynomial-time adversary $\mathcal{A}$ without k whose committed trace has at least q embedding violations,*

$$\Pr[\text{PP-Mark}_{\mathrm{accept}}(I^*, \mathrm{md}^*, \pi^*) = 1 \ \wedge\ (I^*, \pi^*) \notin \mathcal{G}]$$
$$\leq\ \mathrm{Adv}_{\mathsf{H}}^{\mathrm{hash}}(\mathcal{A}) + \mathrm{Adv}_{\mathrm{SP1}}^{\mathrm{sound}}(\mathcal{A}) + \binom{N-q}{k_{\mathrm{open}}} \Big/ \binom{N}{k_{\mathrm{open}}},$$

where $\mathcal{G}$ contains honest image-proof pairs for the target claim, $\mathrm{Adv}^{\mathrm{hash}}$ aggregates the required hash failures, and $\mathrm{Adv}^{\mathrm{sound}}$ is the assumed knowledge-soundness error. The last term is at most $(1 - q/N)^{k_{\mathrm{open}}}$; the bound is per fixed candidate.

Proof sketch; full reduction in Appendix A. Acceptance requires both $\mathrm{score} \geq \tau$ and a receipt matching the expected public statement. Under the assumed knowledge-soundness, a valid receipt yields a witness satisfying $\mathcal{P}_1$–$\mathcal{P}_4$, except with error $\mathrm{Adv}_{\mathrm{SP1}}^{\mathrm{sound}}$.

Binding and commitment. For the pinned claim, recovering a valid secret-key witness or changing a committed trace without changing its root requires a preimage or collision failure, included in $\mathrm{Adv}_{\mathsf{H}}^{\mathrm{hash}}$.

Partial opening. Under the fixed-candidate and uniform-sampling conditions, K misses the q violating positions with probability $\binom{N-q}{k_{\mathrm{open}}} / \binom{N}{k_{\mathrm{open}}}$. Otherwise $\mathcal{P}_4$ rejects. The image-derived challenge alone does not establish these sampling conditions for adaptively selected candidates.

Image binding. A receipt for different canonical pixels fails the public-statement comparison unless their hashes collide. This argument does not require different challenge seeds to produce different subsets K.

Summing the three failure modes yields the stated bound. □

Parameter choice and remaining guarantees. At N=1000, k_{open}=32, and q=125, the exact miss probability is 0.01297 (98.7% spot-check detection); at q=500 it is at most 2^{-32}. The bound weakens for small q, and clean-FPR calibration alone does not bound adversarial FAR. Appendix A gives the scope and fixed-point details.

4 Experiments

Experimental Setup. We use Stable Diffusion v2.1 base with 50 DDIM generation steps and guidance scale 7.5. Generation steps are distinct from inversion steps: the transformation evaluation uses FP32 DDIM-100 inversion and alignment (App. K). PP-Mark uses $\alpha = 4.0$, $N = 1000$, and SHA-256 sampling. Prompts are drawn from the Stable-Diffusion-Prompts dataset; prompt–seed lists and frozen calibration settings accompany the reproduction materials. All methods use native detection statistics calibrated to 1% FPR on clean SD2.1 images. SDXL results are in App. G; BRISQUE, KID, and CLIP comparisons are in App. L (Table 20).

Runtime. PP-Mark$_{\mathrm{score}}$ retains the submitted 0.89 s embedding and 0.72 s verification times. SP1 v6.3.1 at $N = 1000$, reports 17.0 s projected receipt issuance and a 2.81 MB artifact. Tables 8–10 distinguish the timing boundaries and submitted versus updated configurations.

4.1 Verification Ablation

To validate the necessity of zero-knowledge verification (Table 2), we compare four verification architectures under two targeted attacks on $n = 100$ watermarked images at FPR = 1%:

Table 3: Forgery Acceptance Rate (FAR) under two attacks across four verification architectures ($n = 100$, FPR = 1%).

Attack	Sig-only	Sig+Score	PP-Mark (score)	PP-Mark (accept)
Forged binding	1.00	1.00	$\leq 0.01^{\dagger}$	**0.00**
Compliance bypass	1.00	0.00	$\leq 0.01^{\dagger}$	**0.00**

†Upper-bounded by FPR; 0/100 observed. Binding b is cryptographically fixed by the secret key.

Forged binding. An adversary who controls only the metadata signing key $\mathrm{sk}_{\mathrm{sig}}$, a standard digital-signature key used by the Sig-only and Sig+Score baselines to authenticate metadata, independent of and unrelated to PP-Mark's binding key k, fabricates a binding b' to maximize the detection score for a target image. We evaluate 10,000 random bindings per image (sufficient to expect $\sim$70 scores exceeding τ under the null per-trial probability of $\approx 0.7\%$) and report the fraction for which the best score exceeds τ.

Compliance bypass. An adversary generates a clean (unwatermarked) image and signs legitimate-looking metadata without embedding any watermark signal.

Table 3 reports the results. Sig-only fails both attacks; Sig+Score blocks compliance bypass but not forged binding because random search over 10,000 bindings consistently finds one whose induced sign pattern exceeds τ (mean best score $3.88 \gg \tau = 2.47$). PP-Mark(score) resists both because the binding is cryptographically fixed by the secret key k; PP-Mark(accept) additionally rejects any fabricated binding via ZKP, confirming Theorem 3.3 empirically.

For exact canonical-image integrity under a trusted signer, an image-hash signature is simpler and sufficient; App. B.1 isolates PP-Mark's added embedding-compliance evidence and transform-tolerant Score.

4.2 Imprint Forgery Attack

Imprint forgery is a strong black-box and method-agnostic threat: from a single watermarked reference image, an attacker attempts to transfer the watermark imprint onto an arbitrary cover so the verifier attributes it to the same source [Müller et al., 2025]. This captures a realistic adversary who does not know the target watermark design but can adapt a proxy generator, making it a practical strong attack. We use SD2.1 for both the target and the proxy in the imprint-forgery attack. To evaluate this threat fairly, we construct a fixed prompt pool (200 prompts; 100/100 calibration/evaluation) and prepare a cover pool of 100 images as targets for attribution transfer. We fix 100 cover–reference pairs (one reference per cover) and reuse the same pairs across all methods to avoid pairing bias.

Direct sanity check. We first apply the standard imprint-forgery optimizer (n=100 per step, $T \in \{50, 100, 150\}$) to Tree-Ring and Gaussian Shading [Yang et al., 2024] to confirm the attack pipeline functions correctly; both are highly vulnerable (TR FAR@1%FPR: 0.92–0.97; GS: 1.00; full breakdown in Table 11, App. F).

Transfer (cross-method). We extend the pipeline cross-method: imprints extracted from a TR- or GS-watermarked reference are optimized onto each fixed cover, with all detectors evaluated at thresholds calibrated to 1% FPR on clean images. Across step budgets $T \in \{50, 100, 150\}$, PP-Mark$_{\mathrm{score}}$ remains at 0.00–0.01, while TrustMark, InvisMark, and other baselines (RingID, StableSig, HiDDeN) exceed the 1% calibration floor (Table 12; bit-accuracy calibration for binary-only detectors in App. F).

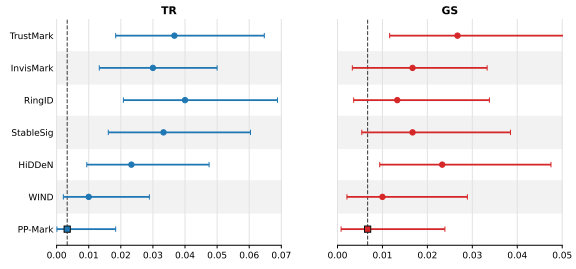

Figure 3: Cross-method transfer at FPR=1%. Error bars denote 95% CI; the vertical line indicates PP-Mark$_{\mathrm{score}}$.

Table 4: Transfer summary over steps. FAR@1%FPR (targeted verify) with 95% CI. All methods evaluated at FPR=1% (clean-calibrated). $\Delta\%$ is computed as $(k - k_{\text{PP-Mark}_{\text{score}}})/k_{\text{PP-Mark}_{\text{score}}} \times 100$, where $k_{\text{PP-Mark}_{\text{score}}} = 1$ (TR) and $k_{\text{PP-Mark}_{\text{score}}} = 2$ (GS). Fisher p is from a one-sided Fisher exact test vs PP-Mark$_{\text{score}}$ (H1: method $>$ PP-Mark$_{\text{score}}$).

	TR				GS			
Method	**FAR@1%FPR**	**95% CI**	**Δ%**	**Fisher p**	**FAR@1%FPR**	**95% CI**	**Δ%**	**Fisher p**
TrustMark	0.0367	0.0184–0.0647	+1000.0	0.0029	0.0267	0.0116–0.0519	+300.0	0.0532
InvisMark	0.0300	0.0138–0.0562	+800.0	0.0102	0.0167	0.0054–0.0385	+150.0	0.2252
RingID	0.0400	0.0208–0.0688	+1100.0	0.0016	0.0133	0.0036–0.0338	+100.0	0.3430
StableSig	0.0333	0.0161–0.0604	+900.0	0.0055	0.0167	0.0054–0.0385	+150.0	0.2252
HiDDeN	0.0233	0.0094–0.0475	+600.0	0.0343	0.0233	0.0094–0.0475	+250.0	0.0882
WIND	0.0100	0.0021–0.0289	+200.0	0.3119	0.0100	0.0021–0.0289	+50.0	0.5000
PP-Mark$_{\text{score}}$	0.0033	0.0001–0.0184	–	–	0.0067	0.0008–0.0239	–	–
PP-Mark$_{\text{accept}}$	0.0000	–	–	–	0.0000	–	–	–

Summary. Table 4 summarizes aggregated transfer acceptance. PP-Mark$_{\text{score}}$ remains at or below the 1% calibration floor (TR: 0.0033, GS: 0.0067), so the imprint-forgery attack fails to move PP-Mark beyond the clean-image noise floor; baselines exhibit substantially higher TR-sourced transfer.

The $\Delta\%$ column is a descriptive summary computed against PP-Mark$_{\text{score}}$'s small denominators; the absolute FARs and 95% CIs provide the primary basis for comparison and consistently rank PP-Mark$_{\text{score}}$ as the least susceptible method.

One-sided Fisher exact tests against PP-Mark$_{\text{score}}$ confirm this ordering: TR-sourced differences for RingID, Stable Signature, and HiDDeN are significant ($p = 0.0016, 0.0055, 0.0343$); GS-sourced differences are not significant at α=0.05, but the point estimates and CIs preserve the same ordering with PP-Mark$_{\text{score}}$ as least vulnerable.

4.3 White-Box Forgery Attack

We consider a strongest case white-box adversary with full access to each detector and gradients, optimizing inputs to maximize the detection score under an ℓ_∞ budget $\epsilon = 2/255$ using $T = 150$ PGD steps. For each method, forgery acceptance is defined by score $\geq \tau$, where τ is calibrated at 1% FPR.

Table 5 shows all baselines reach FAR@1%FPR=1.00, while PP-Mark$_{\text{score}}$ remains lower at 0.76. The smaller mean step than WIND despite lower FAR reflects PP-Mark's keyed image-conditioned objective (variable per-image progress) versus WIND's shared fixed-pattern objective (uniform but later acceptance); cumulative FAR trajectories in Fig. 4.

All methods are attacked via their differentiable pipelines: TrustMark, InvisMark, Stable Signature, and HiDDeN use differentiable decoders; RingID and PP-Mark$_{\text{score}}$ use differentiable inversion-based scores. For WIND, we optimize a differentiable Fourier-distance surrogate while evaluating against the actual detector.

For PP-Mark$_{\text{accept}}$, watermark acceptance requires proof verification. The attacked image changes the image hash h_{img}, which changes challenge_seed and thus the opening positions and the recomputed opening_digest.

Table 5: White-Box forgery outcomes across methods. Avg. Steps to FA reports, among forgery acceptances, the mean optimization step at which score $\geq \tau$ first occurs. PSNR is averaged against the original cover image.

Method	**FAR @1%FPR**	**Avg. Steps to FA**	**Avg. PSNR ↑**
TrustMark	1.00	2.41	75.73
InvisMark	1.00	3.07	72.67
StableSig	1.00	5.40	65.91
HiDDeN	1.00	8.08	62.96
RingID	1.00	2.26	74.51
WIND	1.00	22.20	55.76
PP-Mark$_{\text{score}}$	0.76	17.42	67.69
PP-Mark$_{\text{accept}}$	0.00	–	–

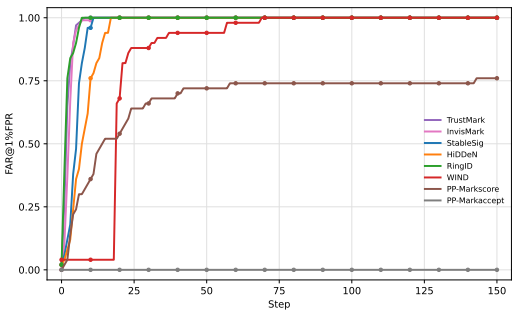

Figure 4: Cumulative FAR@1%FPR vs. step (same attack budget across all methods).

As a result, the verifier's opening check fails and the proof verification rejects, keeping PP-Mark$_{\mathrm{accept}}$ at FAR=0.

4.4 Removal Attacks

PP-Mark survives standard watermark erasure across three attack families at FPR$\leq 1\%$. **Regeneration** (diffusion re-synthesis [Zhao et al., 2024]): PP-Mark TPR 0.98 vs StableSig/HiDDeN 0.01–0.06, InvisMark ≤ 0.05, TrustMark 0.85$\rightarrow$0.39 across noise steps; latent-space embedding survives where pixel-space methods are washed out. **Steganalysis** (K reference samples): PP-Mark TPR 0.864–0.963 vs Tree-Ring 0.02–0.038, StableSig 0.17–0.20. **Pixel/geometric transformations**: JPEG/blur/noise/brightness 0.86–1.00; random crop SD 2.1: 0.87, SDXL: 0.74; rotation 75° via rotation-sync (SD 2.1: 0.85, SDXL: 0.97). Full configurations and per-attack results are in Apps. H, I, and K (Table 19).

5 Conclusion

PP-Mark combines context-bound statistical detection with public proof verification of a committed embedding relation. Score supports transformation-tolerant screening; Accept additionally checks an image-bound receipt under the stated trust and sampling assumptions. Experiments on SD2.1 and SDXL demonstrate the resulting forgery resistance in the evaluated settings. The proof checks embedding consistency rather than re-executing the generator. Limitations (proof overhead, key lifecycle, geometric robustness, and generative-model scope) are discussed in App. M.

6 Broader Impacts

This work provides a decentralized, verifiable mechanism for AI image provenance, helping reduce misattribution and synthetic media misuse. By removing reliance on centralized APIs and addressing forgery threats, PP-Mark enables transparent auditing of provenance claims under clear operational assumptions. At the same time, stronger provenance and watermark-verification tools may increase the traceability of generated media in some deployment contexts, which could be undesirable where anonymity or sensitive expression is important. We therefore view PP-Mark as a supporting provenance tool that should be deployed with appropriate transparency and user-facing safeguards.

Acknowledgements and Disclosures

This research was supported by the Institute for Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Ministry of Science and ICT (MSIT), Republic of Korea, through the Graduate School of Metaverse Convergence Program (RS-2022-00156318); by the Ministry of Culture, Sports and Tourism (MCST), Republic of Korea, through the Korea Creative Content Agency (KOCCA) under the Culture, Sports and Tourism R&D Program (RS-2023-00219237 and RS-2026-25506280); and by the National Research Foundation of Korea (NRF) grant funded by the Ministry of Science and ICT (MSIT), Republic of Korea (RS-2026-25481217). This research was also supported by the NVIDIA Academic Grant Program using NVIDIA DGX Spark.

References

Kasra Arabi, Benjamin Feuer, R. Teal Witter, Chinmay Hegde, and Niv Cohen. Hidden in the noise: Two-stage robust watermarking for images. In *International Conference on Learning Representations (ICLR)*. ICLR, 2025. doi: 10.48550/arXiv.2412.04653.

Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security (CCS)*, pages 62–73. ACM, 1993. doi: 10.1145/168588.168596.

Mikołaj Bińkowski, Danica J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD GANs. In *International Conference on Learning Representations (ICLR)*, 2018.

Tu Bui, Shruti Agarwal, and John Collomosse. TrustMark: Robust watermarking and watermark removal for arbitrary resolution images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 18629–18639, October 2025.

Robert Chesney and Danielle K. Citron. Deep fakes: A looming challenge for privacy, democracy, and national security. *California Law Review*, 107(6):1753–1819, 2019. doi: 10.15779/Z38RV0D15J.

Hai Ci, Pei Yang, Yiren Song, and Mike Zheng Shou. RingID: Rethinking tree-ring watermarking for enhanced multi-key identification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 338–354, 2024. doi: 10.1007/978-3-031-73390-1_20.

Coalition for Content Provenance and Authenticity (C2PA). Content credentials: C2PA technical specification. Technical report, C2PA, 2026. URL `https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html`. Version 2.4. Accessed: 2026-05-03.

Pierre Fernandez, Guillaume Couairon, Hervé Jégou, Matthijs Douze, and Teddy Furon. The Stable Signature: Rooting watermarks in latent diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22466–22477, 2023. doi: 10.1109/ICCV51070.2023.02053.

Junfeng Guo, Yiming Li, Ruibo Chen, Yihan Wu, Chenxi Liu, and Heng Huang. ZeroMark: Towards dataset ownership verification without disclosing watermark. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-3829.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830.

Zhengyuan Jiang, Jinghuai Zhang, and Neil Zhenqiang Gong. Evading watermark based detection of AI-generated content. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23)*, pages 1168–1181, Copenhagen, Denmark, November 2023. ACM. doi: 10.1145/3576915.3623189.

Sung Ju Lee and Nam Ik Cho. Semantic watermarking reinvented: Enhancing robustness and generation quality with fourier integrity. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 18759–18769, Paris, France, October 2025. IEEE/CVF.

Liangqi Lei, Keke Gai, Jing Yu, and Liehuang Zhu. DiffuseTrace: A transparent and flexible watermarking scheme for latent diffusion model. *arXiv preprint arXiv:2405.02696*, 2024. doi: 10.48550/arXiv.2405.02696.

Junhua Lin and Marc Juarez. A crack in the bark: Leveraging public knowledge to remove Tree-Ring watermarks. In *34th USENIX Security Symposium (USENIX Security 25)*, pages 7331–7348, Seattle, WA, August 2025. USENIX Association. ISBN 978-1-939133-52-6. URL `https://www.usenix.org/conference/usenixsecurity25/presentation/lin-junhua`.

Shayne Longpre, Robert Mahari, Naana Obeng-Marnu, William Brannon, Tobin South, Katy Ilonka Gero, Alex Pentland, and Jad Kabbara. Data authenticity, consent, & provenance for AI are all broken: what will it take to fix them? In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 32711–32725. PMLR, 21–27 Jul 2024.

Wenhao Luo, Zhangyi Shen, Ye Yao, Feng Ding, Guopu Zhu, and Weizhi Meng. TraceMark-LDM: Authenticatable watermarking for latent diffusion models via binary-guided rearrangement. *Expert Systems with Applications*, 316:131831, 2026. doi: 10.1016/j.eswa.2026.131831.

Po-Yuan Mao, Cheng-Chang Tsai, and Chun-Shien Lu. MaXsive: High-capacity and robust training-free generative image watermarking in diffusion models. In *Proceedings of the 33rd ACM International Conference on Multimedia (MM '25)*, pages 11443–11452, Dublin, Ireland, October 2025. ACM. doi: 10.1145/3746027.3755266.

Ralph C. Merkle. A digital signature based on a conventional encryption function. In Carl Pomerance, editor, *Advances in Cryptology – CRYPTO '87*, volume 293 of *Lecture Notes in Computer Science*, pages 369–378. Springer, 1988. doi: 10.1007/3-540-48184-2_32.

Andreas Müller, Denis Lukovnikov, Jonas Thietke, Asja Fischer, and Erwin Quiring. Black-box forgery attacks on semantic watermarks for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20937–20946, June 2025. doi: 10.1109/CVPR52734.2025.01950.

Uma Roy. Introducing SP1: A performant, 100% open-source, contributor-friendly zkvm, February 2024. URL `https://blog.succinct.xyz/introducing-sp1/`. Accessed: 2026-04-24.

Kareem Shehata, Aashish Kolluri, and Prateek Saxena. CLUE-Mark: Watermarking diffusion models using CLWE. *arXiv preprint arXiv:2411.11434*, 2024. doi: 10.48550/arXiv.2411.11434.

James C. Spall. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control*, 37(3):332–341, 1992. doi: 10.1109/9.119632.

Matthew Tancik, Ben Mildenhall, and Ren Ng. StegaStamp: Invisible hyperlinks in physical photographs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2117–2126, June 2020. doi: 10.1109/CVPR42600.2020.00219.

Ke Wang, Zishuo Zhao, Xinyuan Song, Zelin Li, Libin Xia, Chris Tong, Bill Shi, Wenjie Qu, Eric Yang, and Lynn Ai. VeriLLM: A lightweight framework for publicly verifiable decentralized inference. *arXiv preprint arXiv:2509.24257*, 2025. doi: 10.48550/arXiv.2509.24257. Version 4, revised 2026-01-22.

Yuxin Wen, John Kirchenbauer, Jonas Geiping, and Tom Goldstein. Tree-Rings watermarks: Invisible fingerprints for diffusion images. In *Advances in Neural Information Processing Systems*, volume 36, pages 58047–58063, 2023. doi: 10.52202/075280-2529.

Rui Xu, Mengya Hu, Deren Lei, Yaxi Li, David Lowe, Alex Gorevski, Mingyu Wang, Emily Ching, and Alex Deng. InvisMark: Invisible and robust watermarking for AI-generated image provenance. In *Proceedings of the Winter Conference on Applications of Computer Vision (WACV)*, pages 909–918, February 2025. doi: 10.1109/WACV61041.2025.00098.

Zijin Yang, Kai Zeng, Kejiang Chen, Han Fang, Weiming Zhang, and Nenghai Yu. Gaussian Shading: Provable performance-lossless image watermarking for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12162–12171, 2024. doi: 10.1109/CVPR52733.2024.01156.

Lijun Zhang, Xiao Liu, Antoni Viros Martin, Cindy Xiong Bearfield, Yuriy Brun, and Hui Guan. Attack-resilient image watermarking using stable diffusion. In *Advances in Neural Information Processing Systems*, volume 37, pages 38480–38507, 2024. doi: 10.52202/079017-1215.

Xuandong Zhao, Kexun Zhang, Zihao Su, Saastha Vasan, Ilya Grishchenko, Christopher Kruegel, Giovanni Vigna, Yu-Xiang Wang, and Lei Li. Invisible image watermarks are provably removable using generative AI. In *Advances in Neural Information Processing Systems*, volume 37, pages 8643–8672, 2024. doi: 10.52202/079017-0276.

Jiren Zhu, Russell Kaplan, Justin Johnson, and Li Fei-Fei. HiDDeN: Hiding data with deep networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 682–697, 2018. doi: 10.1007/978-3-030-01267-0_40.

Yifan Zhu, Yihan Wang, and Xiao-Shan Gao. Towards robust content watermarking against removal and forgery attacks. *arXiv preprint arXiv:2604.06662*, 2026. doi: 10.48550/arXiv.2604.06662. CVPR 2026 Findings.

A Full Proof of Theorem 3.3

We give the step-by-step reduction for Theorem 3.3 (conditions on $\mathcal{A}$ and q are as stated in the main text).

Security game. Fix the verifier-pinned producer claim (b, h_{ctx}) and let $\mathcal{G}$ contain its honest (I, π) pairs. The adversary receives the verification key, public artifacts, and hash-oracle access, but not k. The verifier reconstructs the expected public statement, including the canonical image hash, and rejects mismatching proof public values. We consider one candidate image and trace fixed before its first opening-challenge query, as required by Theorem 3.3; the adversary wins if this candidate is accepted outside $\mathcal{G}$.

Witness extraction. Under the knowledge-soundness assumption in Assumption 3.2, an accepting receipt yields a witness $w^* = (k^*, c^*, \mathcal{T}^*, \mathrm{paths}^*)$ satisfying $\mathcal{P}_1$–$\mathcal{P}_4$, except with error $\mathrm{Adv}^{\mathrm{sound}}_{\mathrm{SP1}}$. The following argument conditions on successful extraction.

Case A: Binding forgery. The verifier fixes (b, h_{ctx}) independently of the candidate. By $\mathcal{P}_1$, the extracted key satisfies $b = \mathsf{H}(\texttt{ctx_hash} \| h_{\mathrm{ctx}} \| k^*)$. Recovering such a secret-key witness requires a preimage failure; a distinct key for the pinned input additionally yields a second preimage. These events are included in the aggregate hash-failure term $\mathrm{Adv}^{\mathrm{hash}}_{\mathsf{H}}$.

Case B: Merkle trace forgery. $\mathcal{P}_2$ requires each leaf $\mathrm{leaf}_i = \mathsf{H}(i, g_i, z_i^{(0)}, \texttt{bit}_i)$ to authenticate to the committed root [Merkle, 1988]. Changing any tuple after fixing that root requires a hash collision, included in $\mathrm{Adv}^{\mathrm{hash}}_{\mathsf{H}}$. A newly committed trace is covered by the partial-opening case below.

Case C: Partial opening with q cheating positions. Let $F \subseteq S$ contain the q positions violating $\mathcal{P}_4$ in Sec. 3.8: the encoded payload bit, normalized LUT baseline, or Q18 embedding equality is incorrect. Since $\mathcal{P}_4$ checks all three conditions on each opened index, an accepting witness must have $K \cap F = \emptyset$.

The opening set is deterministic in the challenge seed,

$$\mathrm{challenge_seed}^* \;=\; \mathsf{H}\big(\texttt{ppmark_opening_v1} \,\|\, h^*_{\mathrm{img}} \,\|\, h^*_{\mathrm{ctx}} \,\|\, \mathrm{sample_root}^*\big)\,,$$

and $K = \mathrm{DeriveOpen}(\mathrm{challenge_seed}^*, k_{\mathrm{open}}, N)$. The candidate's canonical image, trace, and protocol parameters are fixed before this first challenge query. This is an explicit condition of the bound, not a consequence of hashing alone.

Modelling H as a random oracle [Bellare and Rogaway, 1993] (standard for SHA-256 in this regime), conditionally on all commitments being fixed before the seed is evaluated, K is a uniformly random size-k_{open} subset of $[N]$. The probability that a uniform k_{open}-subset avoids the q-element set F is

$$\Pr[K \cap F = \emptyset] = \frac{\binom{N-q}{k_{\mathrm{open}}}}{\binom{N}{k_{\mathrm{open}}}} = \prod_{j=0}^{k_{\mathrm{open}}-1} \frac{N-q-j}{N-j}.$$

Since $(N-q-j)/(N-j) \le (N-q)/N = 1 - q/N$ for every $j \in \{0, \ldots, k_{\mathrm{open}} - 1\}$ (the ratio is non-increasing in j, attaining its maximum $1 - q/N$ at $j = 0$), the telescoping product yields

$$\Pr[K \cap F = \emptyset] \;\le\; (1 - q/N)^{k_{\mathrm{open}}} \;\le\; \exp(-k_{\mathrm{open}} \cdot q/N)\,.$$

Conditioned on $K \cap F \neq \emptyset$, $\mathcal{P}_4$ rejects the witness with probability 1, hence π^* fails verification. The contribution of Case C is therefore $(1 - q/N)^{k_{\mathrm{open}}}$.

Case D: Image binding (cross-image proof transfer). Consider two images with distinct canonical pixel arrays. If their hashes differ, an existing receipt fails comparison with the target image's expected public statement. If their hashes coincide, this is a hash collision, included in $\mathrm{Adv}^{\mathrm{hash}}_{\mathsf{H}}$. Different seeds need not produce different subsets K; receipt binding relies on the checked public statement. Files with identical canonical pixels represent the same artifact for this guarantee.

Union bound. Combining the failure probabilities from witness extraction and Cases A–D,

$$\Pr[\mathcal{A}\text{ wins}] \;\le\; \mathrm{Adv}^{\mathrm{sound}}_{\mathrm{SP1}}(\mathcal{A}) + \mathrm{Adv}^{\mathrm{hash}}_{\mathsf{H}}(\mathcal{A}) + \frac{\binom{N-q}{k_{\mathrm{open}}}}{\binom{N}{k_{\mathrm{open}}}},$$

matching the conditional, per-candidate statement of Theorem 3.3. □

Numerical instantiation. At $N = 1000$, $k_{\mathrm{open}} = 32$, and $q = 125$, the exact without-replacement miss probability is 0.012965 (approximately 98.7% detection). The simpler upper bound $(1-q/N)^{32}$ is 0.013940. This upper bound is approximately $2^{-13.28}$ at $q = 250$ and 2^{-32} at $q = 500$. These are partial-opening probabilities under the theorem's sampling conditions; cryptographic failure terms remain separate.

Small-q regime ($q < 0.125N$). The upper bound $(1 - q/N)^{32}$ ranges from approximately 0.014 to 0.968 for $q \in \{1, \ldots, 125\}$, so partial opening alone gives a weak guarantee for small q. Score checks all N positions, but clean-FPR calibration does not establish a general adversarial FAR bound. The submitted per-step transfer FARs remain 0–0.01 (Sec. 4.2); the separate white-box FAR results remain as reported in Table 5. These empirical results do not prove a small-q security bound.

Opening-tolerance parameter ε. We use $\varepsilon = 0$: $\mathcal{P}_4$ independently reconstructs the normalized Gaussian baseline and payload sign, then checks exact equality under the declared Q18 arithmetic. A nonzero tolerance changes the verified relation and would require a separate security analysis.

B Verification Design Details

This appendix complements Sec. 3.7 with the design rationale for using a cryptographic hash rather than a perceptual hash for h_{img}, and Sec. 4.1 with image-hash signature controls.

Cryptographic hash vs perceptual hash for proof binding. We use SHA-256 over canonical pixels for h_{img} rather than a perceptual hash. The verifier compares the receipt's public statement with the expected image hash, so a change to canonical pixels invalidates an existing receipt except for a hash collision. A perceptual identifier shared by distinct images would weaken this exact-image binding.

Benign edits that change canonical pixels therefore invalidate the original Accept receipt. Routine Score screening of a derivative requires no re-proof; a fresh receipt is needed only to grant Accept to that derivative itself. 17.0 s projected receipt issuance at $N = 1000$ (Table 10). Exact integrity concerns canonical pixels rather than the original file-container bytes.

B.1 Image-hash signature controls

The Sig-only and Sig+Score baselines in Sec. 4.1 authenticate metadata. We additionally evaluate two domain-separated Ed25519 controls with a fixed provider public key and the same canonical-pixel hash h_{img}: **Bare-Sig** signs h_{img}; **Context-Sig** signs $h_{\mathrm{img}} \parallel h_{\mathrm{ctx}} \parallel b$.

Table 6: Image-hash signature controls (n=100 per test).

Test	Bare-Sig	Context-Sig	PP-Mark evidence
Original verifies	100/100	100/100	—
Cross-image replay	0/100	0/100	—
Context substitution	100/100	0/100	—
Authorized signer signs an unwatermarked image	100/100	100/100	0/100 (Score/Accept)
JPEG (q=25) after signing	0/100	0/100	94/100 (Score)
75% crop–resize after signing	0/100	0/100	87/100 (Score)

Signature columns count valid signatures. PP-Mark entries are the separately reported compliance-bypass results (Sec. 4.1) and Score detections (Table 19).

For exact canonical-image integrity under a trusted signer, an image-hash signature has the same modification-rejection role as Accept's image-hash binding and is simpler and preferable. Context-Sig also prevents context substitution, but an authorized signer can still sign an image without embedding a watermark. A signature authenticates the signer's assertion; it does not verify that embedding occurred. Accept additionally requires both a passing watermark score and a valid receipt enforcing $\mathcal{P}_1$–$\mathcal{P}_4$, including sampled embedding consistency. Score provides transform-tolerant in-image detection.

C Empirical Verification of Assumption 3.1

We verify the two conditions of Assumption 3.1 on $n = 100$ watermarked images by computing the DDIM inversion residual $\epsilon = \hat{z}_0 - z_0$ (channel 0) for each image, where z_0 is the watermarked latent saved during generation and $\hat{z}_0$ is the latent recovered via DDIM inversion from the decoded pixel image. We report results on both SD 2.1 (latent 64×64) and SDXL (latent 128×128).

Zero-mean condition. On SD 2.1 the per-image residual mean is $\bar{\epsilon} = +0.011 \pm 0.003$ (grand mean $\pm$ SE across 100 images), two orders of magnitude below the residual standard deviation $\sigma_\epsilon \approx 0.50$, so $|\bar{\epsilon}|/\sigma_\epsilon \approx 0.02$. On SDXL the grand mean is $\bar{\epsilon} = -0.022 \pm 0.004$ (95% CI $[-0.031, -0.013]$) with $\sigma_\epsilon \approx 0.44$, yielding $|\bar{\epsilon}|/\sigma_\epsilon \approx 0.05$. Each channel-0 residual contains 4,096 entries on SD 2.1 and 16,384 on SDXL. The bias is statistically detectable across the 100 images but small relative to both the residual standard deviation and α_{eff}; its magnitude alone does not establish a zero effect on the score.

Bounded correlation condition. On SD 2.1 the mean Pearson correlation is $|\rho(\epsilon, s)| = 0.30$ (median 0.30, max 0.48). On SDXL the correlation is larger: $|\rho| = 0.54$ (median 0.55, max 0.66). This two-model comparison does not isolate latent dimensionality as the cause of the difference. Per-image correlations are negative in every case in both models, so the attenuation acts as a multiplicative shrinkage on the score rather than as random noise. The per-image p-values are small (all < 0.05 on SDXL). We therefore report correlation magnitudes and observed score margins; correlation below one alone is not a sufficient detection guarantee. On SD 2.1 the watermark signal contribution $\alpha_{\text{eff}}\sqrt{N} = 30.7$ is an expectation reference, not a hard score maximum. The observed mean is 25.7 (approximately 16% lower), and the minimum score is 17.4, about $7.0\times$ the threshold $\tau = 2.47$ across the 100 images. On SDXL the observed mean clean score is 13.6 ($\text{score}/\tau \approx 5.5$), again with every clean image crossing τ at FPR$\leq$1%. These measurements support the residual approximation on the evaluated images; both models retain 100/100 detections despite their different score margins.

Figure 5 overlays the SD 2.1 and SDXL distributions on shared axes: (a) per-image residual means are centered near zero for both models ($\bar{\epsilon} = +0.011$ on SD 2.1, -0.022 on SDXL), and (b) per-image $|\rho(\epsilon, s)|$ concentrates around 0.30 for SD 2.1 and shifts to approximately 0.54 for SDXL while remaining bounded well below 1.

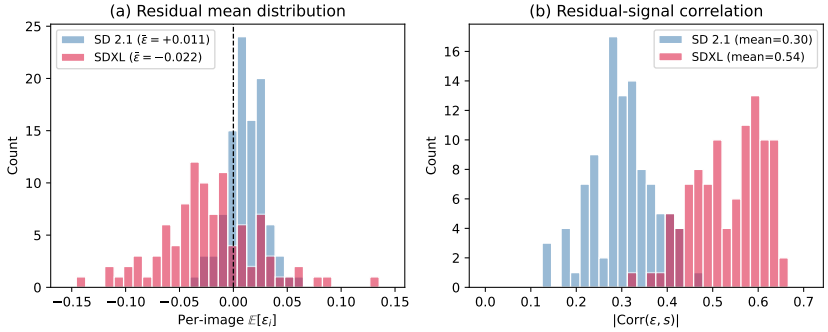

Figure 5: Empirical assessment of Assumption 3.1 on SD 2.1 and SDXL (n=100 each, overlaid). (a) Per-image residual means cluster near zero. (b) Mean $|\text{Corr}(\epsilon, s)|$ is 0.30 on SD 2.1 and 0.54 on SDXL; the two-model comparison does not isolate the cause of this difference.

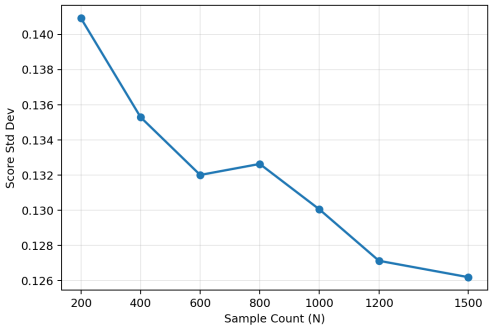

(a) Score concentration: standard deviation vs. sample count N.

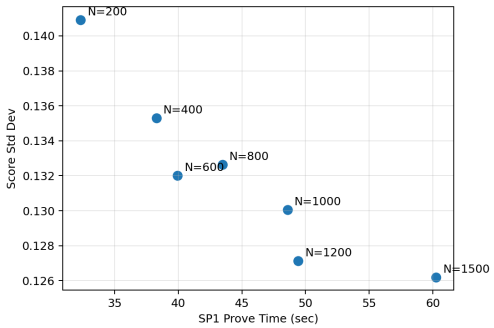

(b) Knee trade-off between proving cost and score variability.

Figure 6: Sample size ablation. (a) Score variance decreases as N grows; (b) knee point identifies $N{=}1000$ as the trade-off between proving cost and score stability.

D Ablation Study

D.1 Sample Size and Statistical Reliability

We analyze the trade-off between statistical reliability and computational overhead by varying the number of sampled latent positions $N \in \{200, \ldots, 1500\}$. We identify $N = 1000$ as the reasonable point where statistical stability is maximized without incurring excessive prover latency or bandwidth costs.

Statistical Concentration: Theoretically, the detection score (cosine similarity) acts as a smooth function of bounded aggregates, suggesting Hoeffding-sense concentration. Empirically, the score's standard deviation across images decreases by 7.7% as N increases from 200 to 1000 ($0.1409 \rightarrow 0.1301$). However, increasing N further to 1500 yields only a marginal 3.0% additional reduction, indicating diminishing returns in variance reduction.

Table 7: SP1 timing vs. sample count N. Submitted v5 measurements are retained; the final row shows the latest configuration with explicitly marked time estimates.

N	Prove / issue (s)	Verify (s)
200	32.309	4.468
400	38.278	5.640
600	39.932	5.886
800	43.465	6.651
1000	48.572	7.589
1200	49.425	8.877
1500	60.272	9.802
$1000^{\dagger}$	$17.0^{\dagger}$	$4.0^{\dagger}$

†SP1 v6.3.1, latest circuit implementation. The 17.0 s, 4.0 s values are a projected fresh-process receipt-issuance phase total, including setup, execution, proving, self-verification, and serialization.

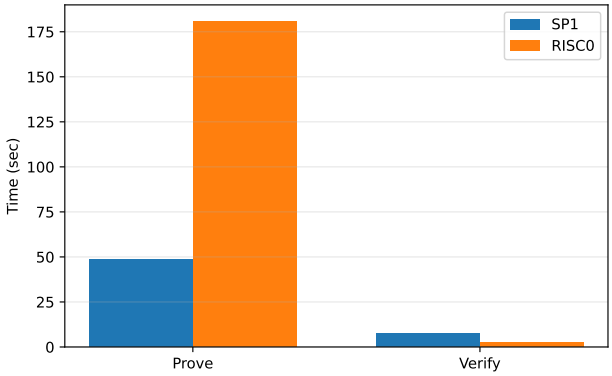

Figure 7: Submitted prover and verifier timings at N=1000 (α=4.0). SP1 uses a single H100 GPU for proving with CPU receipt verification; RISC0 uses the same SD2.1 setup and public inputs with CPU receipt verification.

Computational Cost: In the submitted SP1 v5 configuration, prover and verification times increase with N (Table 7): proving rises from 32.3 s ($N = 200$) to 48.6 s ($N = 1000$). Figure 7 retains the submitted SP1/RISC0 comparison. The appended v6.3.1 row uses a different implementation and timing boundary and is not part of the original N sweep.

D.2 Runtime Comparison

Table 8 retains the submitted SD2.1 embedding and detection timings and updates the Accept row to the latest rebuttal configuration. Semantic methods embed during diffusion generation; post-hoc methods apply a separate encoder. Updated receipt issuance is approximately 17.0 s projected at $N = 1000$. Table 7 gives a conditional estimate of standalone receipt verification under unchanged fixed overhead. Total Accept-verification time has not been jointly measured and is left unreported.

Table 8: Runtime comparison on SD2.1. Submitted method timings are retained; the Accept row uses the latest rebuttal receipt-issuance value.

Method	Embed (s)	Verify (s)
Tree-Ring	0.89	0.75
Gaussian Shading	3.10	0.72
RingID	0.89	0.72
WIND	0.89	1.10
Stable Signature	0.91	0.008
HiDDeN	0.001	0.001
TrustMark	0.013	0.004
InvisMark	0.002	0.009
PP-Mark (score)	0.89	0.72
PP-Mark (accept, v6.3.1)	$0.89 + 17.0^{\dagger}$	$0.72 + 4.0^{\dagger}$

†The 17.0 s 4.0 s are projected receipt issuance, including initialization and self-verification, rather than standalone proving. The displayed sum combines the submitted embedding time with that projection; it is not a jointly timed run. Total Accept verification is not jointly measured; the separate receipt-verification estimate is in Table 7.

D.3 SP1 Receipt Size

Table 9 retains the submitted size sweep and appends the latest SP1 v6.3.1 Core artifact (2.81 MB at $N = 1000$). Table 10 reproduces the final rebuttal's backend comparison. The post-submission implementation uses an optimized full-trace streaming commitment; the original Merkle configuration remains the basis of the submitted sweep. Version, implementation, and timing-boundary differences preclude attributing the change to the SP1 release alone.

Table 9: SP1 receipt byte sizes across N (SD 2.1, α=4.0). Submitted rows are retained, followed by the latest rebuttal configuration.

N	Public inputs (B)	Witness (KB)	Receipt (MB)
300	674	254.69	12.24
400	664	339.21	16.87
600	665	553.15	24.81
1000	666	921.60	40.68
1500	666	1493.57	65.96
1000†	1,004	117.0	2.81

†Latest SP1 v6.3.1 Core implementation, updated public-input, witness, receipt sizes.

Table 10: Proof-backend comparison from the final rebuttal ($N = 1000$), including the latest circuit implementation.

Proof configuration	Time (s)	Artifact
SP1 v6.3.1 †	17.0†	2.81 MB
SP1 v5	48.572	40.68 MB
RISC0	180.9	40.51 MB
PLONK	312.677	3,874 B
Groth16	438.380	1,499 B

†SP1 v6.3.1 uses a Core receipt. Its 17.0 s value is the projected fresh-process issuance phase total; the 2.81 MB artifact size is measured. The other rows retain the times and sizes reported.

D.4 Watermark Strength

We study the effect of watermark strength α on visual quality while fixing the sampling budget at $N = 1000$. We report BRISQUE (lower is better) and focus on the 95th percentile across images to capture tail behavior. For each $\alpha \in \{3.0, 3.5, 4.0, 4.5\}$ we evaluate 100 watermarked images under identical settings.

The BRISQUE p95 is relatively stable from $\alpha = 3.0$ to 4.0 (roughly 34.1–34.9), but it jumps at $\alpha = 4.5$ to 40.5, indicating a heavier quality tail. This suggests that pushing α beyond 4.0 yields a disproportionate increase in worst-case artifacts without a clear compensating gain in reliability, given that $N = 1000$ already stabilizes detection. We therefore select $\alpha = 4.0$ as a conservative balance between robustness and perceptual quality (Fig. 8).

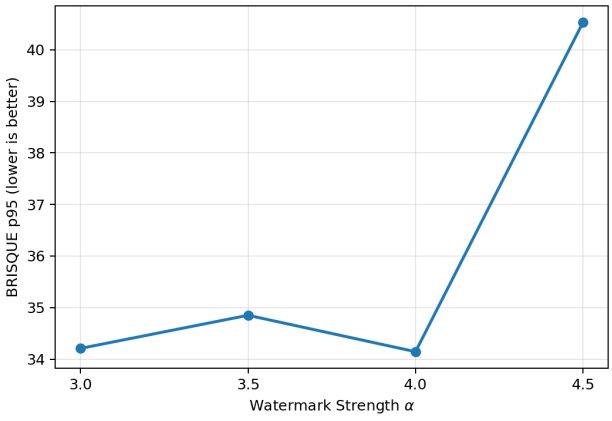

Figure 8: BRISQUE p95 vs. watermark strength α (lower is better).

E Computational Resources

The submitted experiments were conducted on an HPE server equipped with a single NVIDIA H100 PCIe GPU (80 GB) and an Intel® Xeon® Gold 6542Y CPU (24 cores, 48 threads, 2.9 GHz, 60 MB cache). The GPU handles image generation, watermark embedding, latent inversion, attack optimization (PGD, SPSA, imprint forgery, regeneration), and SP1 STARK proving; SP1 receipt verification runs on the CPU. Per-method runtimes are in Appendix D.2, and the submitted N-dependent prove/verify sweep is in Appendix D.1 (Table 7). Post-submission configurations are distinguished in Table 10.

F Imprint Forgery Details

This appendix complements Sec. 4.2 with the full step-wise breakdowns for both the Direct and Transfer imprint-forgery experiments, along with the bit-accuracy calibration used for binary-only detectors (TrustMark, InvisMark). The aggregated cross-method comparison (FAR with 95% confidence intervals, $\Delta\%$ ratios, and one-sided Fisher exact tests) and the corresponding forest plot remain in the main paper (Table 4, Figure 3).

F.1 Direct Imprint Forgery

We sanity-check the standard imprint-forgery optimizer on Tree-Ring (TR) and Gaussian Shading (GS) themselves to confirm the attack pipeline functions correctly before applying it cross-method. For each fixed cover–reference pair, we run the optimizer for $T \in \{50, 100, 150\}$ steps and report the Forgery Acceptance Rate at 1% FPR (FAR@1%FPR, n=100 per step) using each method's native detector with thresholds calibrated to 1% FPR on clean (non-watermarked) images.

Table 11: Direct vulnerability under targeted imprint forgery. FAR@1%FPR (targeted verify); all methods evaluated at FPR=1% (clean-calibrated).

	TR	**GS**
Step	**FAR@1%FPR**	**FAR@1%FPR**
50	0.92	1.00
100	0.97	1.00
150	0.97	1.00
Avg.	0.953	1.00

Both methods are highly vulnerable across all step budgets. Tree-Ring's FAR rises from 0.92 (50 steps) to 0.97 (100–150 steps), averaging 0.953 across the three step counts; Gaussian Shading saturates at 1.00 across all step counts, indicating near-perfect imprint transfer onto arbitrary covers. These rates confirm the optimizer reliably extracts and re-injects watermark imprints, providing a sound baseline for cross-method evaluation in App. F.2.

F.2 Transfer Imprint Forgery

We extend the imprint-forgery pipeline cross-method: imprints extracted from a TR- or GS-watermarked reference are optimized onto each fixed cover and evaluated against each target method's native detector. All thresholds are calibrated to 1% FPR on clean images, so the reported differences reflect detector susceptibility under identical pairings, step budgets, and calibration. Step budget is $T \in \{50, 100, 150\}$ with n=100 pairs per step.

Bit-accuracy calibration for binary-only detectors. TrustMark [Bui et al., 2025] and InvisMark [Xu et al., 2025] expose only binary detection flags through their public APIs; standard threshold-based FPR calibration is therefore not directly available. We instead use bit accuracy, the fraction of correctly decoded payload bits, as a continuous score, enabling threshold-based FPR=1% calibration on clean images. The resulting calibration thresholds are τ=0.58 for TrustMark (clean bit accuracy $\approx$0.48, watermarked $\approx$1.00) and τ=0.49 for InvisMark (clean $\approx$0.41, watermarked $\approx$0.96). The

large clean–watermarked gaps confirm bit accuracy provides a reliable continuous score for these methods despite their binary-only public interfaces.

Table 12: Transfer results by method and step. FAR@1%FPR (targeted verify); all methods evaluated at FPR=1% (clean-calibrated). TrustMark and InvisMark use bit accuracy as the continuous score.

Source	Step	TrustMark	InvisMark	RingID	StableSig	HiDDeN	WIND	PP-Mark$_{\text{score}}$	PP-Mark$_{\text{accept}}$
TR	50	0.05	0.06	0.01	0.04	0.04	0.01	0.00	0.00
	100	0.03	0.01	0.10	0.04	0.02	0.01	0.00	0.00
	150	0.03	0.02	0.01	0.02	0.01	0.01	0.01	0.00
GS	50	0.01	0.03	0.01	0.01	0.01	0.01	0.00	0.00
	100	0.02	0.01	0.02	0.02	0.03	0.01	0.01	0.00
	150	0.05	0.01	0.01	0.02	0.03	0.01	0.01	0.00

Across all step budgets and both imprint sources (TR, GS), PP-Mark$_{\text{score}}$ remains at 0.00–0.01, indicating that the standard imprint-forgery pipeline rarely pushes its detection score beyond the 1% calibration floor. The peak baseline value is RingID at 0.10 under TR-sourced imprints with 100 optimization steps; TrustMark, InvisMark, StableSig, and HiDDeN range across 0.01–0.06 across the various source–step combinations. Among baselines, WIND remains the closest competitor to PP-Mark$_{\text{score}}$ with FAR consistently at 0.01, while PP-Mark$_{\text{accept}}$ stays at 0.00 across all settings, confirming that the cryptographic verification rejects any imprint-altered image regardless of score-domain optimization.

G Evaluation on SDXL

We further evaluate PP-Mark on Stable Diffusion XL (SDXL 1.0), a second latent-diffusion model sharing the initial-latent embedding and DDIM-inversion interface with Stable Diffusion 2.1 (SD2.1). SDXL generates 1024×1024 images with a 128×128 latent space (four times the dimensionality of SD2.1's 64×64 latent). Among the baselines from the main experiments, only TrustMark and InvisMark are natively compatible with SDXL's resolution; the remaining methods (RingID, StableSig, HiDDeN, WIND) either predate SDXL or require resolution downsampling that renders their outputs non-comparable. We therefore scope this evaluation to PP-Mark, TrustMark, and InvisMark.

Setup. We generate 100 clean and 100 watermarked SDXL images per method (50 DDIM steps). PP-Mark uses α=4.0, sample count N=1000 with the 128×128 latent. All thresholds are calibrated at FPR$\leq$1% on SDXL clean images (n=100).

Detection. All three methods achieve TPR=1.00 on their respective SDXL watermarked images (n=100 each). PP-Mark's watermarked scores (mean=13.63) exceed τ=2.47 by $5.5\times$, confirming that the larger latent space preserves detection margin.

Direct imprint vulnerability (TR/GS sanity check). As in SD2.1 (Table 11), we first verify that the SDXL imprint-forgery pipeline reaches each semantic watermark's own detector. Table 13 reports FAR@1%FPR for Tree-Ring and Gaussian Shading under direct attack (n=100 per step). Both are highly vulnerable (TR: 0.83–0.91, GS: 1.00 across steps), confirming the SDXL attack pipeline is functional before evaluating pixel-space and post-hoc methods below.

Transfer imprint forgery. We apply the same imprint-forgery protocol as Section 4.2 using Tree-Ring (TR) and Gaussian Shading (GS) as imprint sources (n=100 per step). Table 14 reports FAR at the calibrated thresholds.

PP-Mark$_{\text{score}}$ achieves FAR=0.00–0.02, consistent with SD2.1 (0.00–0.01). PP-Mark$_{\text{accept}}$ remains at 0.00 across all conditions, as the ZKP proof verification deterministically rejects forged images. TrustMark remains robust (FAR$\leq$0.05). InvisMark shows materially higher vulnerability on SDXL (FAR=0.09–0.17) compared to SD2.1 (0.01–0.06), suggesting that SDXL's larger latent space amplifies imprint transfer susceptibility for pixel-space decoders. Figure 9 visualizes pooled FAR

Table 13: Direct vulnerability under targeted imprint forgery on SDXL (sanity check). FAR@1%FPR (targeted verify); TR and GS evaluated at their own clean-calibrated thresholds.

	TR	GS
STEP	FAR@1%FPR	FAR@1%FPR
50	0.83	1.00
100	0.90	1.00
150	0.91	1.00
AVG.	0.880	1.00

with Clopper–Pearson 95% confidence intervals: PP-Mark$_{\text{score}}$ 0.01 [0.00, 0.03], TrustMark 0.03 [0.01, 0.05], InvisMark 0.13 [0.09, 0.17] for TR (GS similar).

Table 14: Transfer imprint forgery on SDXL. FAR at FPR$\leq$1% thresholds calibrated on SDXL clean (n=100 per step).

		TR	GS
Method	**Step**	**FAR@1%FPR**	**FAR@1%FPR**
TrustMark	50	0.03	0.00
	100	0.05	0.02
	150	0.00	0.01
InvisMark	50	0.16	0.15
	100	0.09	0.11
	150	0.14	0.17
PP-Mark$_{\text{score}}$	50	0.02	0.01
	100	0.00	0.01
	150	0.01	0.01
PP-Mark$_{\text{accept}}$	50	0.00	0.00
	100	0.00	0.00
	150	0.00	0.00

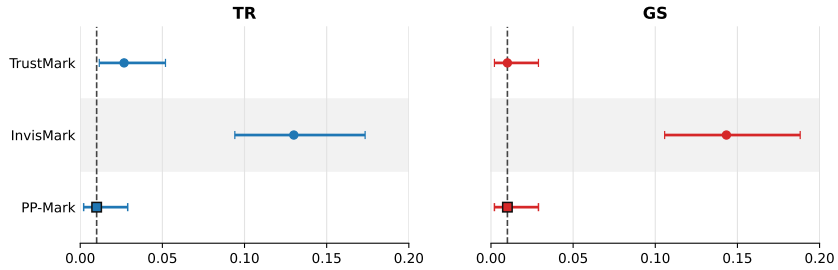

Figure 9: SDXL transfer imprint forgery: pooled FAR across 3 steps (n=300 per source). Error bars denote Clopper–Pearson 95% CI; the dashed line marks PP-Mark$_{\text{score}}$.

Table 15: White-box PGD forgery on SDXL (ϵ=2/255, T=150, n=100). FAR at FPR$\leq$1% under any-step acceptance; Avg. PSNR averaged against the original cover image at the submitted-image step (first-breach step for acceptances, final step otherwise).

Method	**FAR**	**Avg. Steps to FA**	**Avg. PSNR ↑**
TrustMark	1.00	2.78	80.28
InvisMark	1.00	3.83	71.22
PP-Mark$_{\text{score}}$	0.11	29.55	85.63
PP-Mark$_{\text{accept}}$	0.00	–	–

White-box PGD forgery. We apply the same PGD attack (ϵ=2/255, T=150) to all three methods on SDXL, using the any-step acceptance criterion (attack succeeds if score$\geq\tau$ at any step) as in SD2.1. Table 15 reports outcomes and Figure 10 plots the cumulative FAR trajectories.

TrustMark and InvisMark are breached at 100% within 3–4 steps on average, consistent with SD2.1. PP-Mark$_{\text{score}}$ achieves FAR=0.11 (n=100, 95% Clopper–Pearson CI $[5.62\%, 18.83\%]$), a substantial improvement over SD2.1 (FAR=0.76). The 4$\times$ larger latent dimensionality (65,536 vs 16,384) may contribute to this difference, but the two-model comparison does not isolate its effect. The models also differ in PSNR: PP-Mark$_{\text{score}}$'s average rises from 67.69 (SD2.1) to 85.63 (SDXL), whereas TrustMark/InvisMark PSNR moves only slightly (+4.5/$-$1.5 dB). Although the attacker fills the full ℓ_∞ budget (ϵ=2/255), the gradient in the larger SDXL latent is spatially sparse, so most pixels remain unperturbed and the effective MSE stays small, consistent with the lower FAR. PP-Mark$_{\text{accept}}$ remains at 0.00 by construction.

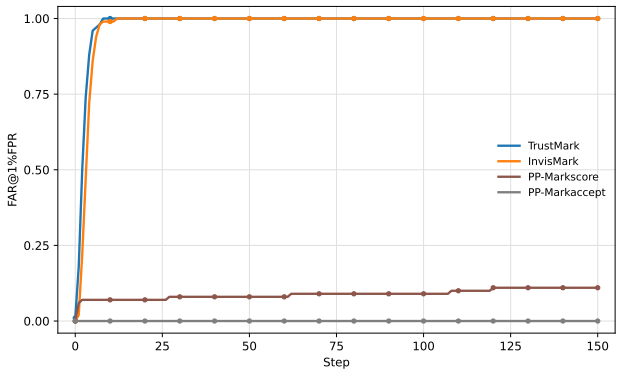

Figure 10: SDXL cumulative FAR@1%FPR vs. optimization step under white-box PGD (ϵ=2/255, T=150, n=100).

Summary. On the evaluated SDXL samples, accept mode maintains FAR=0.00 under both imprint transfer and white-box PGD attacks, and score mode has white-box FAR=0.11 compared with 0.76 on SD2.1. These results cover two models sharing a latent-diffusion/DDIM-inversion interface; generalization to substantially different generative architectures remains untested.

Table 16: Post-Regeneration attack detection rate at FPR=1% thresholds; each noise step uses 100 images; noise step is 30/60/100.

	NOISE STEP		
METHOD	30STEPS	60STEPS	100STEPS
STABLESIG	0.04	0.02	0.03
HIDDEN	0.10	0.06	0.04
TRUSTMARK	0.85	0.63	0.39
INVISMARK	0.01	0.05	0.02
TREE-RING	0.95	0.93	0.93
PP-MARK	0.98	0.98	0.98

H Regeneration Attack

We evaluate detection rates after regeneration at each noise level to measure how well the watermark survives re-synthesis. Table 16 reports detection rates after regeneration.

Table 17: Steganalysis removal summary (TPR@1%FPR) by K.

	K			
Method	**10**	**100**	**500**	**1000**
StableSig	0.200	0.200	0.173	0.171
Tree-Ring	0.190	0.020	0.038	0.035
TrustMark	0.990	0.960	0.960	0.960
InvisMark	1.000	1.000	1.000	1.000
WIND	1.000	1.000	1.000	1.000
PP-Mark	0.864	0.959	0.963	0.958

I Steganalysis Attack

We sweep K={10, 100, 500, 1000} and evaluate the post-attack detection rate at the calibrated FPR=1% threshold. Table 17 shows that StableSignature and Tree-Ring detection drops sharply as K increases, indicating increasing removal effectiveness. In contrast, TrustMark, InvisMark, WIND, and PP-Mark remain highly detectable even at large K, because their watermark patterns are content-adaptive rather than fixed, making the averaged residual a poor estimate of any individual image's watermark. Notably, for all content-adaptive methods the averaged residual converges to zero as K grows, confirming that steganalysis is fundamentally ineffective against this encoder class. The minor TPR variation across these methods (e.g. PP-Mark 0.96 vs. InvisMark 1.00 at $K \geq 100$) reflects differences in score-to-threshold margin rather than actual robustness: PP-Mark's tighter calibration leaves less headroom for the residual noise added at small K, whereas InvisMark's wider margin absorbs it entirely.

J SPSA-based Black-Box Attacks

We evaluate PP-Mark under gradient-free SPSA [Spall, 1992] attacks in both removal and forgery settings.

Removal. Given a watermarked image and its binding metadata, the adversary minimizes the verifier's score via SPSA over 50 steps. The score drops below the calibrated threshold τ early, but visual quality collapses immediately (PSNR ≈ 4.7–6.9). Under any reasonable quality constraint (PSNR ≥ 30), no valid outputs are obtained. Figure 11 shows the score and PSNR traces.

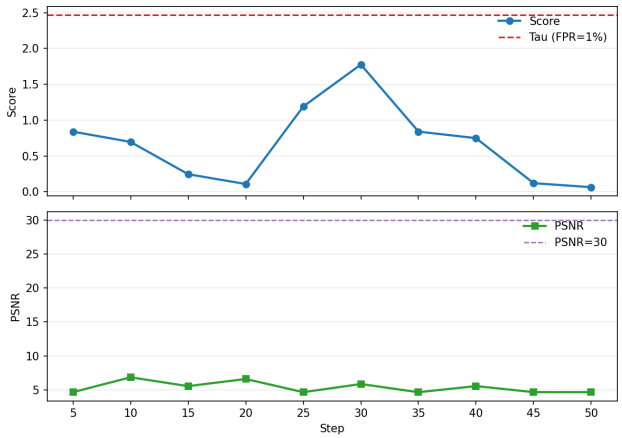

Figure 11: SPSA removal trace. Top: score vs step with τ (FPR=1%); bottom: PSNR vs step.

Table 18: SPSA-based forgery attack on PP-Mark.

Step	**FAR@1%FPR**	**PSNR↑**
50	0.02	6.55
100	0.00	6.61
150	0.01	6.63
Avg.	0.01	6.60

Forgery. The adversary maximizes the verifier's score against a target binding via SPSA, using DDIM inversion to initialize the latent. Table 18 summarizes the results: targeted FAR remains at 0.01 on average, and the optimization merely introduces destructive distortion (PSNR $\approx$ 6.6).

Table 19: Post-transformation detection rate (TPR at FPR$\leq$1%). Individual transformations and the sequential pipeline are evaluated in score mode.

Transform	SD 2.1	SDXL
JPEG compression (q=25)	0.94	1.00
Gaussian blur (8×8)	1.00	1.00
Additive noise (σ=0.1)	0.86	0.95
Brightness jitter (×0.6)	0.98	1.00
Random crop + resize (75%)	0.87	0.74
Rotation 75°	0.85	0.97
75% crop → resize → JPEG (q=25)	0.80†	—

†Success requires detection after both crop–resize and subsequent JPEG (80/100). We use the same 100 SD 2.1 images, bicubic resizing, no survivor filtering, and the submitted FP32 DDIM-100 alignment with a frozen threshold. —: not evaluated.

K Image Transformations

We evaluate watermark robustness under individual and sequential image transformations and report the *score-mode* detection rate (TPR at FPR$\leq$1%). Table 19 summarizes results.

Pixel-level transforms. PP-Mark maintains high detection under lossy compression, blur, noise, and brightness changes. Even aggressive JPEG compression (q=25) preserves 94% detection, and Gaussian blur achieves perfect retention.

Crop and resize. Random cropping to 75% of the original area followed by resizing to the original dimensions is handled by a two-stage alignment combining integer-shift search and sub-pixel refine-

ment. This recovers 87% detection on SD 2.1 (64×64 latent) and 0.74 on SDXL (128×128 latent, $n=50$); the VAE maps overlapping spatial regions to correlated latent coefficients, and the refinement grid covers the residual sub-pixel misalignment.

Rotation. Large rotations are recovered through a dedicated rotation-sync search over $\pm90°$. Rotation at $75°$ retains 0.85 detection on SD 2.1 and 0.97 on SDXL; the larger 128×128 latent of SDXL provides additional redundancy that favours sync recovery. This single angle is reported as a representative worst-case large rotation; the same procedure applies uniformly across angles.

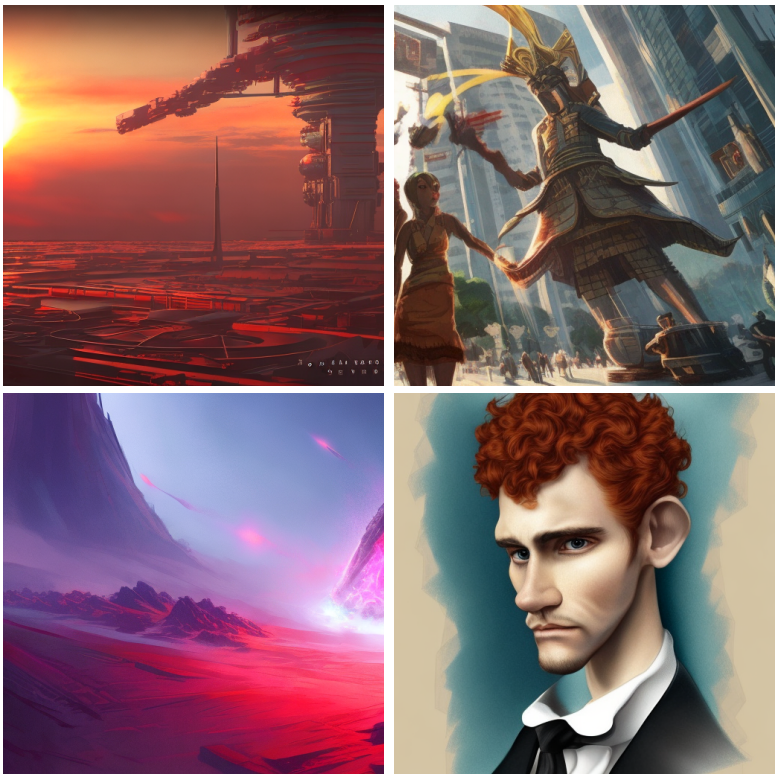

Figure 12: Examples of PP-Mark watermarked images

L Image Quality and Distributional Fidelity

Table 20 reports BRISQUE p95 for no-reference image quality, KID [Bińkowski et al., 2018] for distributional fidelity, and mean CLIP image–prompt cosine similarity (ViT-B/32) for semantic alignment. We use KID rather than FID because the 2,048-dimensional FID covariance is rank-deficient at this sample size.

Table 20: Image quality comparison. KID is mean $\pm$ SD; arrows indicate preferred directions.

Method	BRISQUE p95 $\downarrow$	KID $\times 10^3$ $\downarrow$	CLIP $\times 100$ $\uparrow$
PP-Mark[†]	34.14	1.101 ± 0.202	31.64
Gaussian Shading	27.35	0.621 ± 1.096	32.42
HiDDeN	105.86	7.239 ± 1.474	31.30
InvisMark	31.36	-0.900 ± 0.895	32.53
RingID	34.64	-0.989 ± 0.701	32.41
Stable Signature	32.69	-0.814 ± 0.878	32.50
Tree-Ring	41.90	-0.090 ± 0.922	32.30
WIND	27.42	-0.386 ± 0.879	32.28
TrustMark	29.44	-0.174 ± 1.001	32.54

[†]PP-Mark retains the submitted BRISQUE (n=100, 25 prompts $\times$ 4 seeds) and KID (n=1000, subset 500, 20 splits). Baseline BRISQUE/KID and all CLIP results use the common 100 prompt–seed pairs. Baseline KID uses a disjoint pool of 900 clean reference images; its SD is across 1,000 draws of 100 reference images, with the 100 method images fixed.

Small negative unbiased KID estimates can occur at finite sample sizes and do not denote negative distances. PP-Mark's submitted KID remains low, but the differing sampling protocols do not support fine cross-method KID rankings. PP-Mark does not attain the lowest BRISQUE or highest CLIP similarity; this quality profile accompanies the forgery-resistance results in Sec. 4.2.

M Limitations and Discussions

Generative-model scope. Our experiments cover SD2.1 and SDXL, which share a latent-diffusion/DDIM-inversion interface. The proof circuit verifies the committed embedding trace without re-simulating the generator, while end-to-end score verification requires a compatible initial-latent embedding and recovery interface. Generalization to substantially different generative architectures remains untested.

ZKP Overhead. The latest rebuttal implementation (SP1 v6.3.1, $N = 1000$) produces a 2.81 MB Core receipt and reports approximately 17.0 s projected fresh-process receipt issuance (Table 10). The original N sweeps are retained in Tables 7 and 9; Table 7 also gives a conditional, unmeasured estimate of standalone receipt verification. This is tolerable for archival or on-request verification but costly for per-upload streaming pipelines; a STARK-to-SNARK wrapper or proof aggregation are natural reductions that we leave to future work.

Random Crop on SDXL. Across geometric transformations, random crop with resize on SDXL is the single transformation where detection drops meaningfully (0.74 on SDXL vs. 0.87 on SD 2.1 under the same two-stage alignment; see Table 19). The 128×128 latent of SDXL, four times the dimensionality of SD 2.1, amplifies sub-pixel misalignment introduced by cropping and resampling, so residual near-miss scores cluster just below τ. Refining the sub-pixel alignment grid or incorporating a learned alignment model that undoes VAE-crop interactions would close this remaining gap; we leave this to future work.

Reliance on Explicit Proof Artifacts. Traditional diffusion watermarks delegate trust to a proprietary verifier; PP-Mark removes this centralized dependency to enable public auditing, but in exchange requires an explicit proof artifact to accompany the image. If this artifact is detached, deleted, or corrupted, the image reverts to an unverified state, a challenge inherent to certificate-based provenance frameworks. PP-Mark$_{\text{accept}}$ is designed to certify bit-exact originals; it correctly rejects modified images, limiting its applicability to scenarios where the original is preserved (e.g., legal evidence, regulatory audit). For general distribution environments where images undergo lossy transforms, PP-Mark$_{\text{score}}$ provides statistical verification via soft correlation over the recovered latent, retaining high detection after common transforms (e.g., JPEG q=25: 94%, blur: 100%; see Table 19). Future work should explore anchoring proof artifacts in content-addressed decentralized storage and on-chain registries to ensure persistence independent of any single server.

Key Lifecycle Operational Guidance. PP-Mark's unforgeability is contingent on the producer's binding key k remaining secret; if k leaks, the adversary can mint valid receipts under the victim's identity. We treat this as an operational security requirement and recommend the following lifecycle, consistent with standard PKI practice: (i) *rotation*, each producer maintains a small set of active binding keys $\{k^{(t)}\}$ indexed by a rotation epoch t; metadata includes the current key identifier, so verifiers can tell which epoch a receipt belongs to; (ii) *revocation*, a compromised $k^{(t)}$ is published on a public revocation list (analogous to CRL/OCSP), and verifiers reject receipts bound to revoked epochs; (iii) *forward binding*, because b only depends on the *context* at generation time, rotated keys do not invalidate legitimate historical receipts as long as the corresponding epoch remains non-revoked. Integration with C2PA Content Credentials is straightforward: the epoch ID, rotation policy, and revocation URL can be carried as standard assertion fields.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: Contributions are stated in the abstract and Introduction. The conditional security scope is specified in Section 3.8 and Appendix A; evaluated settings are reported in Sections 4.1–4.3.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Discussed in Appendix M.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Assumption 3.1 is stated in Section 3.6 and empirically assessed in Appendix C. Assumption 3.2 and the fixed-candidate sampling conditions accompany Theorem 3.3 in Section 3.8; Appendix A gives the corresponding proof and scope.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Generation, calibration, and evaluation settings are described in Section 4 and the experimental appendices. The GitHub release will provide code, configurations, and reproduction instructions identifying available inputs, external dependencies, and limitations of the retained historical data.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in

some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: Code was supplied as supplementary material; the updated code and reproduction instructions will be released on GitHub for the camera-ready version.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: PP-Mark uses no learned parameters; calibration and per-attack settings are documented in Section 4 and the appendices.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: FAR confidence intervals use the reported sample counts: individual settings use $n = 100$, while Table 4 pools 100 pairs over three step budgets (300 outcomes per source). Cross-method comparisons there use one-sided Fisher exact tests; Appendix L specifies the sampling protocols and SDs for KID.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.

- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Submitted hardware is described in Appendix E; per-method runtimes are in Appendix D.2. The original SP1 timing sweep is in Appendix D.1 (Table 7); sizes and the rebuttal backend comparison are in Appendix D.3. Updated rows distinguish projected issuance time, estimated receipt-verification time, and measured artifact size, and identify unreported quantities.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: Conforms with the NeurIPS Code of Ethics; no human subjects or sensitive data.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Discussed in Section 6.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: No new generative models, datasets, or scraped data are released.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All baselines and the SP1 zkVM are cited; pretrained Stable Diffusion 2.1 and SDXL 1.0 weights are used under their public Hugging Face licenses.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: The updated PP-Mark codebase, documentation, configurations, and reproduction instructions are being prepared for the camera-ready GitHub release.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: No crowdsourcing or human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

Justification: No human subjects.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: Core methodology does not use LLMs.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.