SPECTRALCANARY: EMBARRASSINGLY SIMPLE LOW-FREQUENCY TRACES FOR T2I DATASET-USE AUDITING

Anonymous authors
Paper under double-blind review

ABSTRACT

Auditing whether a text-to-image (T2I) model was fine-tuned on a particular image collection is difficult: a model can absorb collection-level visual characteristics without reproducing training images. Proactive dataset-use auditing plants a shared signal in a dataset before release, then later checks whether that signal appears in a suspect model's generated outputs. Existing approaches often trade off verifiability, fidelity, robustness, and deployability. We introduce **SpectralCanary**, which uses a smooth low-pass mask to blend a fixed canary image into every collection image while leaving captions unchanged and requiring neither auxiliary generation nor per-image optimization. A lightweight verifier detects the transferred signal in suspect-model outputs. Across eight tests on the Pokemon dataset and three realistic data transformations, spanning two architecture families, SpectralCanary achieves the highest average per-generation TPR@1% FPR (91.1%) and fixed-threshold model-level accuracy using 30 generations (94.8%) among the evaluated methods. Across five datasets spanning artistic, medical, satellite, and fashion domains with SD3.5, SpectralCanary averages 92.4% per-generation TPR@1% FPR while preserving the released images (SSIM 0.989, LPIPS 0.022) with a near-zero mean change FID (mean ΔFID 0.05) across datasets. Matched-energy ablations on Pokemon and ROCOv2 support our design choice, low-frequency placement transfers more reliably compared to mid/high-frequency placement. Under the evaluated adaptive attacks, the most effective removal requires aggressive modification and incurs more than 25 ΔFID points. These results identify low-frequency visual signals as effective carriers for collection-level T2I audit traces.

1 INTRODUCTION

Creative and scientific image collections embody substantial human labor, expertise, and investment. Adapting text-to-image (T2I) models to such collections can support legitimate personalization, education, design, and scientific communication (Rombach et al., 2022; Esser et al., 2024; Hu et al., 2022). The same capability, however, allows a third party to fine-tune and deploy a model on a collection without the rights holder's consent, attribution, or compensation. Because a model can absorb a collection's visual characteristics without reproducing individual images, rights holders may have little evidence that their work was used (Pang & Wang, 2025; Li et al., 2025). Practical dataset-use auditing can strengthen accountability and licensing compliance while allowing images to remain available for legitimate use.

We study proactive dataset-use auditing, commonly termed data traceability or unauthorized data-use verification in prior works. A practical method should satisfy four requirements. **(Req. 1) Collection-level verifiability.** It should provide evidence from black-box outputs even when no training image is reproduced. **(Req. 2) Fidelity and utility preservation.** Protected samples should preserve their original appearance and the utility of models trained on them, avoiding harm to legitimate use while resisting easy filtering. **(Req. 3) Pipeline robustness.** The trace should remain detectable across model architectures, after augmentation such as image regeneration or recaptioning. **(Req. 4) Operational deployability.** The protection transform should be model-agnostic and scale efficiently to large collections without auxiliary generation or iterative per-image optimization.

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

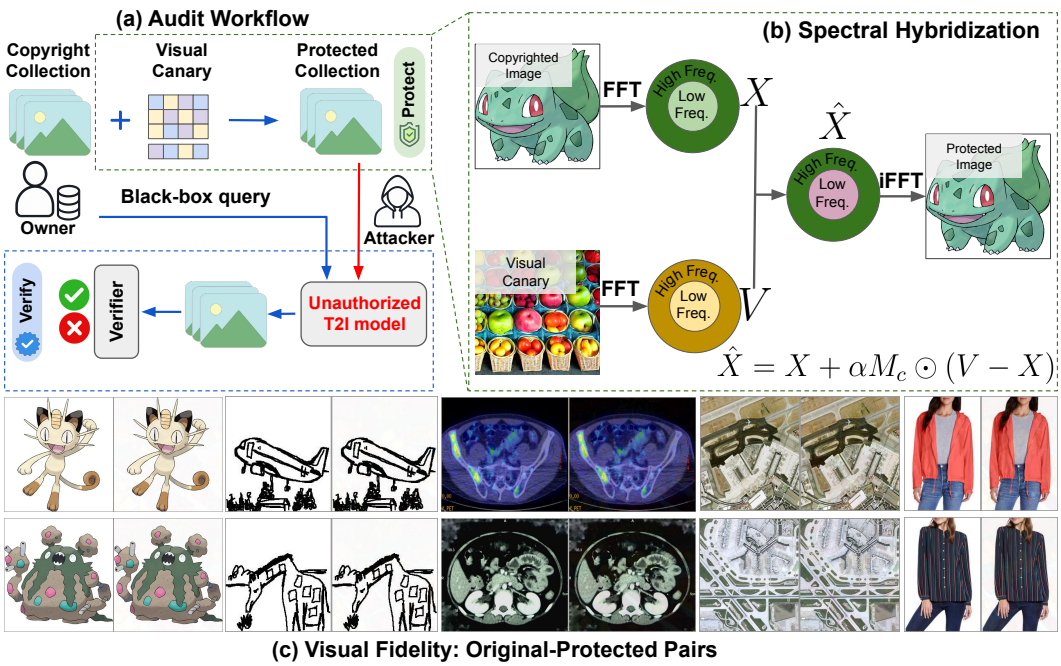

Figure 1: **SpectralCanary overview. (a) Auditing workflow.** An owner protects a collection before release, then queries a suspect T2I model and tests whether its generated outputs contain the canary. **(b) Spectral hybridization.** SpectralCanary blends a fixed canary through a smooth low-pass mask M_c with strength α, retaining most of the remaining spectrum. A lightweight verifier trained on original–protected pairs detects the trace in model generations. **(c) Visual fidelity.** Original–protected pairs across five datasets, including sparse sketches.

Prior approaches typically sacrifice at least one requirement. **Reactive auditing: limited collection-level evidence (Req. 1).** Membership and similarity tests rely on sample-specific evidence that generalization may erase (Pang & Wang, 2025; Li et al., 2025). **Pixel-level watermarks: limited robustness (Req. 3).** Weak signals may be attenuated by augmentation, regeneration, or fail to transfer across model backbones (Wang et al., 2024b; Li et al., 2025). **Explicit canaries: fidelity and robustness trade-offs (Reqs. 2–3).** Repeated templates can be conspicuous and filterable, while recaptioning can remove textual triggers (Ren et al., 2025). **Adjacent data-poisoning objectives.** Poisoning induces chosen behavior rather than provides evidence of dataset use. Persistent behavior may not be unique or reliably elicited, so it does not by itself establish use of a particular collection (Wang et al., 2024a; Jang et al., 2025). Across proactive methods, auxiliary generation, filtering, and iterative optimization can also limit deployability **(Req. 4)** by increasing computational cost or requiring external models, domain-specific assumptions (Ren et al., 2025; Li et al., 2025; Jang et al., 2025). These limitations motivate a subtle, transferable, caption-independent, and inexpensive trace.

We propose **SpectralCanary**, whose key design choice is the frequency band used for the shared trace. Before release, the dataset owner chooses one fixed canary image. For each collection image, SpectralCanary mixes a small amount of the canary's low-frequency Fourier content into the original while retaining most of its spectrum and leaving its caption unchanged (Fig. 1(b)). The protected spectra are thereby shifted toward a common low-frequency target, creating a collection-level signal. Low frequencies encode broad spatial structure that neural networks preferentially learn (Rahaman et al., 2019; Xu et al., 2019) and that is important to T2I generation (Yi et al., 2024; Ma et al., 2026). This motivates our hypothesis that a weak low-frequency trace can transfer during fine-tuning. The dataset owner then queries a suspect model and uses a lightweight verifier to test whether its generated outputs contain the trace (Fig. 1(a)). Importantly, SpectralCanary uses a single deterministic transform requiring no caption changes, auxiliary generator, or per-image optimization.

We evaluate all four requirements. Robustness tests cover a convolutional latent-diffusion U-Net (Rombach et al., 2022), a transformer-based rectified-flow model (Esser et al., 2024); data augmentation, image regeneration (Zhao et al., 2024a), and recaptioning. Direct baselines include pixel-level traces (SIREN (Li et al., 2025), DIAGNOSIS (Wang et al., 2024b)) and explicit canaries

(EnTruth (Ren et al., 2025)). We also adapt Silent Branding (Jang et al., 2025) but report its evaluation separately in Appx. Q, because it targets visual branding rather than dataset-use auditing.

Auditing effectiveness and robustness. For per-generation discrimination, AUC measures verifier-score separation, while TPR at 1% FPR measures detection of outputs from protected-data models. Across eight Pokemon settings, SpectralCanary achieves 98.0% AUC and 91.1% TPR. In fixed-threshold model-level verification using 30 generations, it achieves 94.8% average accuracy under the same settings. **Fidelity and utility preservation.** SSIM and LPIPS compare original and protected images, while ΔFID measures the change in generative quality. Across five visual domains with SD3.5, SpectralCanary achieves 99.5% AUC and 92.4% TPR with SSIM 0.989, LPIPS 0.022, with a near-zero mean change FID (ΔFID 0.05). On sparse Sketch-Scene drawings, it achieves 100% AUC and TPR with SSIM 0.997 and LPIPS 0.002. **Operational deployability.** Protection takes 0.39 seconds per image on our CPU setup (Appx. E). **Mechanism and boundary.** In matched-energy ablations on Pokemon and ROCOv2, we show that low-frequency placement transfers more reliably compared to mid/high-frequency placement. Within the evaluated adaptive attacks, effective removal requires aggressive modification and incurs more than 25 ΔFID points.

Our main contributions are:

- We identify low-frequency visual structure as an effective carrier for weak collection-level traces during T2I fine-tuning. In matched band size and perturbation energy ablations on Pokemon and ROCOv2 with SD1.5, we show that low-frequency placement yields higher per-generation TPR than mid/high-frequency placement.
- We introduce SpectralCanary, a deterministic and caption-preserving spectral hybridization method. It protects each image with one inexpensive transform, requiring neither an auxiliary generator nor per-image optimization.
- We jointly evaluate per-generation discrimination, KS-based multi-query auditing, fixed-threshold model-level verification, image fidelity, generative utility, robustness, and protection cost across two architecture families, five datasets, three direct auditing baselines, an adapted poisoning baseline, and three realistic data transformations. This evaluation establishes a strong empirical operating point and characterizes the boundary under adaptive spectral replacement.

2 SPECTRALCANARY

Problem setup. Let $\mathcal{D} = \{(x_i, t_i)\}_{i=1}^{N}$ be the defender's image–caption collection, with $x_i \in [0,1]^{H\times W\times 3}$. Given a canary $v^\star$, protected fraction $\delta \in (0,1]$, blending strength $\alpha \in [0,1]$, and Gaussian bandwidth $c > 0$, SpectralCanary constructs protected collection $\widehat{\mathcal{D}}$ and trains verifier f_θ (Alg. 1). A third party may fine-tune a T2I model $\mathcal{G}$ on $\widehat{\mathcal{D}}$. The defender audits its black-box generations against clean reference scores (Alg. 2).

Threat model. The algorithm and released collection are public. δ, α, and c may be known. The defender retains the originals, $v^\star$, and f_θ and has only black-box access to $\mathcal{G}$. The audit provides calibrated statistical evidence, not cryptographic attribution or white-box robustness. Revealing the canary spectrum enables the idealized inversion below.

Design motivation. Low frequencies encode broad structure that neural networks preferentially learn (Rahaman et al., 2019; Xu et al., 2019) and that is important to T2I generation (Yi et al., 2024; Ma et al., 2026). They are therefore a plausible trace carrier. However, strong modification is visible. SpectralCanary uses a smooth low-pass mask to weight frequencies and a small blending strength to control the overall change. Our controlled study in Sec. 4.5 shows low-frequency placement yields higher detection rates than mid/high-frequency placement.

Fourier convention. Throughout, $\mathcal{F}_c$ denotes the centered channel-wise 2D discrete Fourier transform and $\mathcal{F}_c^{-1}$ denotes its inverse:

$$\mathcal{F}_c = \mathrm{fftshift} \circ \mathrm{DFT}_2, \qquad \mathcal{F}_c^{-1} = \mathrm{IDFT}_2 \circ \mathrm{ifftshift}\,. \tag{1}$$

We use unitary normalization and apply the transform channel-wise. Any matched forward–inverse scaling gives the same reconstruction because Eq. 4 is homogeneous in X_i and V. Real inputs and the real, conjugate-symmetric mask make the inverse real up to roundoff. Let Π_8 denote clipping to $[0,1]$ followed by 8-bit quantization.

Algorithm 1: Protection and Verifier Training

Input: $\mathcal{D}, v^\star, \delta, c, \alpha$
Output: $\widehat{\mathcal{D}}, f_\theta$

Construct M_c using Eq. 2
$\mathcal{I}_\delta \leftarrow \mathrm{UniformSubset}([N], \lfloor \delta N \rfloor)$
$V \leftarrow \mathcal{F}_c(v^\star)$; $\widehat{\mathcal{D}} \leftarrow \mathcal{D}$
for all $i \in \mathcal{I}_\delta$ **do**
  $X_i \leftarrow \mathcal{F}_c(x_i)$
  $\widehat{X}_i \leftarrow X_i + \alpha M_c \odot (V - X_i)$
  $\tilde{x}_i \leftarrow \Pi_8\Big(\Re[\mathcal{F}_c^{-1}(\widehat{X}_i)]\Big)$
  $\widehat{\mathcal{D}}[i] \leftarrow (\tilde{x}_i, t_i)$
end for
Train f_θ on $\mathcal{I}_\delta$ using Eq. 5
return $\widehat{\mathcal{D}}, f_\theta$

Algorithm 2: Auditing a Suspect Model

Input: $\mathcal{G}$, $\mathcal{P}$, f_θ, clean-data reference scores $\mathcal{R}$, level γ
Output: $\mathcal{S}^{\mathcal{G}}$, p-value $q^{\mathcal{G}}$, decision $a^{\mathcal{G}}$

$\mathcal{S}^{\mathcal{G}} \leftarrow \emptyset$
for all $p_j \in \mathcal{P}$ **do**
  Generate $x_j^{\mathcal{G}} \sim \mathcal{G}(p_j)$
  $s_j \leftarrow f_\theta(x_j^{\mathcal{G}})$
  Add s_j to $\mathcal{S}^{\mathcal{G}}$
end for
$q^{\mathcal{G}} \leftarrow \mathrm{KS}_+(\mathcal{S}^{\mathcal{G}}, \mathcal{R})$
$a^{\mathcal{G}} \leftarrow \mathbb{1}[q^{\mathcal{G}} < \gamma]$
return $\mathcal{S}^{\mathcal{G}}, q^{\mathcal{G}}, a^{\mathcal{G}}$

Stage I: Constructing protected images. Alg. 1 uniformly selects $|\mathcal{I}_\delta| = \lfloor \delta N \rfloor$ images; only these are replaced and used to train the verifier ($\delta = 1$ by default). For $i \in \mathcal{I}_\delta$, we resize $v^\star$ as specified in Appx. A and compute $X_i = \mathcal{F}_c(x_i)$ and $V = \mathcal{F}_c(v^\star)$. We define a smooth radial low-pass mask:

$$M_c(k, \ell) = \exp\left(-\frac{r(k, \ell)^2}{2c^2}\right), \tag{2}$$

where (k, ℓ) indexes the shifted Fourier grid, whose center is $(k_0, \ell_0) = (\lfloor H/2 \rfloor, \lfloor W/2 \rfloor)$, and

$$r(k, \ell) = \sqrt{\left(\frac{k - k_0}{H}\right)^2 + \left(\frac{\ell - \ell_0}{W}\right)^2}. \tag{3}$$

Normalizing the two axes separately also accommodates rectangular images. The mask equals one at the spectral center and decays smoothly with radial distance; a smaller bandwidth c concentrates the weighting more tightly around the lowest frequencies. The protected spectrum is

$$\widehat{X}_i = X_i + \alpha M_c \odot (V - X_i) = (1 - \alpha M_c) \odot X_i + \alpha M_c \odot V, \tag{4}$$

where $\alpha \in [0, 1]$ is the blending strength and $\odot$ denotes elementwise multiplication. We export $\tilde{x}_i = \Pi_8(\Re[\mathcal{F}_c^{-1}(\widehat{X}_i)])$. Eq. 4 moves each weighted coefficient toward target V. It combines shared term $\alpha M_c \odot V$ with image-specific term $(1 - \alpha M_c) \odot X_i$. The collection mean shifts toward V, while variation contracts in low-pass-weighted frequencies. Appx. F reports learning dynamics. Fig. 1(c) and Appx. N show visual fidelity. The deterministic construction requires neither an auxiliary generative model nor iterative per-image optimization. Appx. E reports runtime.

Training the verifier. Let $f_\theta(x) \in (0, 1)$ denote the predicted probability that image x contains the protection signal. Using selected protected images as positives and their original counterparts as negatives, we minimize

$$\mathcal{L}_{\mathrm{ver}}(\theta) = -\frac{1}{|\mathcal{I}_\delta|} \sum_{i \in \mathcal{I}_\delta} \left[\log f_\theta(\tilde{x}_i) + \log(1 - f_\theta(x_i))\right]. \tag{5}$$

Training uses no model generations. Each original–protected pair shares semantic content, making the protection transform the only systematic label-aligned difference. This encourages but does not guarantee trace learning. Transfer to unseen generations is evaluated in Sec. 4.2.

Stage II: Auditing a suspect model. Let $\mathcal{R}$ contain clean-data reference scores and $\mathcal{P}$ be trigger-free audit prompts (collection captions in our experiments). For suspect model $\mathcal{G}$, we collect $\mathcal{S}^{\mathcal{G}}$ and apply a one-sided two-sample Kolmogorov–Smirnov test for larger suspect scores, rejecting at level γ. Experiments use matched clean-data models. Deployment calibrates $\mathcal{R}$ on representative independent models. This yields the KS-based multi-query audit decision. AUC and TPR are per-generation discrimination metrics (Appx. A.1).

Invertibility. If an attacker knows the canary spectrum V, blending strength α, and Gaussian bandwidth c, then M_c is known and Eq. 4 can be inverted as

$$X_i = \frac{\widehat{X}_i - \alpha M_c \odot V}{1 - \alpha M_c}, \tag{6}$$

This recovers pre-export X_i exactly when $\alpha < 1$, but clipping and quantization prevent exact recovery from the released image. Our threat model may reveal α and c but keeps V private. These are threat-model assumptions, not security guarantees.

3 Related Work

Reactive and proactive auditing. Reactive methods test an already trained model using membership inference or similarity between generated and reference images (Chen et al., 2026; Pang & Wang, 2025; Radford et al., 2021; Oquab et al., 2024). Such evidence can weaken when fine-tuning transfers collection-level characteristics without retaining identifiable information about individual samples. Proactive methods instead modify a collection before release. Anti-learning methods aim to prevent or redirect customization (Shan et al., 2023; Zhao et al., 2024b; Van Le et al., 2023; Ahn et al., 2025; Liu et al., 2024). Proactive traceability aims to preserve customization utility while making later use statistically auditable.

Proactive traceability. Pixel-level methods such as DIAGNOSIS (Wang et al., 2024b) and SIREN (Li et al., 2025) insert weak transformations or learned signals that can be detected in generated images. Their subtle signals preserve appearance, but can be attenuated by augmentation, regeneration, or model backbone changes. EnTruth instead inserts explicit canary samples containing a repeated template and textual trigger (Ren et al., 2025). This coherent signal can transfer strongly, but the injected samples are conspicuous and caption triggers can be removed. Our direct comparisons span the two main proactive families including subtle pixel-level traces and explicit visual canary samples.

Relation to poisoning. Data poisoning induces chosen behaviors such as concept corruption, reproduction, or persistent branding (Wang et al., 2024a; Jang et al., 2025). Persistence alone is insufficient for auditing. Auditing instead requires a collection-specific signal detectable under ordinary, trigger-free queries with controlled false positives. SpectralCanary targets this auditing objective rather than model corruption. Appx. S provides further discussion and a method taxonomy.

4 Experiments

4.1 Experimental Setup and How to Read the Results

Our experiments ask four concrete questions. First, can SpectralCanary reveal that a model was trained on a protected collection, even when the model architecture or training data are changed? Second, does adding the trace visibly change the released images or make a model trained on them produce worse images? Third, is protection fast enough for a large collection? Finally, why do low frequencies matter, and what happens when an informed attacker tries to remove the trace?

Datasets and models. We cover distinct image domains: charater art (Pokemon-BLIP-Captions (Pinkney, 2022)), sparse scene drawings (Sketch-Scene (Chowdhury et al., 2022)), medical (ROCOv2 (Rückert et al., 2024)), satellite (RSICD (Lu et al., 2017)), and fashion (FACAD (Yang et al., 2020)). We fine-tune Stable Diffusion v1.5 (SD1.5) (Rombach et al., 2022) and Stable Diffusion 3.5 (SD3.5) (Esser et al., 2024). The main experiments use LoRA, a common parameter-efficient form of T2I adaptation. Appx. G repeats the key test with full-parameter fine-tuning.

Comparisons and stress tests. We compare with SIREN (Li et al., 2025) and DIAGNOSIS (Wang et al., 2024b), which modify training images, and EnTruth (Ren et al., 2025), which adds explicit canary samples. Appx. Q separately adapts Silent Branding (Jang et al., 2025), whose original goal is visual branding rather than dataset-use auditing. On Pokemon, we test each method under ordinary training and after three realistic changes to the training set: image augmentation, diffusion-based recreation of each training image (image regeneration) (Zhao et al., 2024a), and automatic recaptioning with Qwen2.5-VL 7B (Bai et al., 2025).

How to read the metrics. We distinguish the embedded *trace* from two evaluation levels. *Per-generation discrimination* scores each generated image independently and asks whether outputs from protected-data models receive higher verifier scores than outputs from clean-data models. AUC summarizes this score separation over all thresholds. TPR at 1% FPR is the fraction of outputs from protected-data models detected at a threshold that falsely flags at most 1% of outputs from

clean-data models; thus, a TPR of 90% means detecting about 90 of 100 protected-data-model outputs while flagging about 1 of 100 clean-data-model outputs. At the model level, we evaluate two decision protocols. *KS-based multi-query auditing* combines scores from 30 generated images in a statistical test and reports the rejection rate on protected-data models. *Fixed-threshold model-level verification* aggregates 30 generations from each held-out model and applies a threshold fixed after independent calibration, reporting TPR, FPR, precision, F1, and accuracy (Appx. H). For image fidelity, lower LPIPS (Zhang et al., 2018) means that a protected image is closer to its original. For generated-image quality, $\Delta\text{FID} = \text{FID}_{\text{protected}} - \text{FID}_{\text{clean}}$ (Seitzer, 2020); values near or below zero indicate no measured quality loss, whereas large positive values indicate degradation. Unless stated otherwise, SpectralCanary uses $\alpha = 0.02$, $c = 0.20$, and $\delta = 1$. Every model-training and image-sampling setting is repeated with three independent runs. Detailed protocols are in Appx. A.

4.2 Per-Generation Discrimination Survives Changes to Models and Training Data

Method	None		Aug.		Recap.		Regen.		Summary	
TPR@1% FPR ↑	SD1.5	SD3.5	SD1.5	SD3.5	SD1.5	SD3.5	SD1.5	SD3.5	Worst	Avg.
EnTruth	85.3	93.9	76.2	**95.3**	19.4	35.9	88.1	91.7	19.4	73.2
SIREN	66.3	91.9	6.9	18.8	63.6	92.0	44.4	37.2	6.9	52.6
DIAGNOSIS	**99.9**	2.3	**88.3**	8.2	0.4	13.0	43.9	1.7	0.4	32.2
SpectralCanary	97.7	**100.0**	79.2	65.9	**93.6**	**100.0**	**92.5**	**100.0**	**65.9**	**91.1**

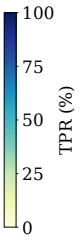

Figure 2: **SpectralCanary maintains strong per-generation discrimination across model and data changes.** Each cell reports per-generation TPR at 1% FPR. "Worst" is the lowest of a method's eight results, and "Avg." is their mean.

SpectralCanary has the strongest average per-generation discrimination across the eight Pokemon settings. Fig. 2 reports per-generation TPR for both model families under ordinary training, training-time augmentation, prompt recaptioning, and image regeneration. SpectralCanary averages 91.1% TPR at 1% FPR, compared with 73.2% for the next-best method. Its weakest individual result is 65.9% under SD3.5 with training-time augmentation; the best competing method has a worst-case result of only 19.4%. No method wins every individual setting: DIAGNOSIS is strongest for ordinary SD1.5 training, and EnTruth is strongest for SD3.5 with training-time augmentation. The complete per-generation AUC, TPR at 5% FPR, and run-level results are in Appx. C. Appx. J separately reports KS-based multi-query auditing across query budgets, showing that more queries generally provide stronger evidence of training on protected data.

Method	None		Aug.		Recap.		Regen.		Summary	
Accuracy ↑	SD1.5	SD3.5	SD1.5	SD3.5	SD1.5	SD3.5	SD1.5	SD3.5	Worst	Avg.
EnTruth	**100.0**	82.0	**100.0**	**90.3**	78.7	50.0	**100.0**	90.0	50.0	86.4
SIREN	**100.0**	**100.0**	52.3	50.0	**100.0**	**100.0**	99.7	50.0	50.0	81.5
DIAGNOSIS	**100.0**	63.7	**100.0**	88.6	50.0	73.1	**100.0**	65.0	50.0	80.1
SpectralCanary	**100.0**	**100.0**	99.6	58.5	**100.0**	**100.0**	**100.0**	**100.0**	**58.5**	**94.8**

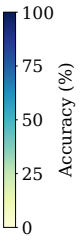

Figure 3: **SpectralCanary is more consistent across transformations and models than the baselines.** Each cell is Acc. "Worst" is the lowest of a method's eight results, and "Avg." is their mean.

Per-generation discrimination translates well to fixed-threshold model-level verification. We aggregate 30 generations per held-out model, following Li et al. (2025), and apply a threshold independently calibrated for each method–architecture pair and fixed across transformations. Fig. 3 reports model-level accuracy across the same eight settings. Each method–architecture–transformation setting includes five clean-data and five protected-data models, yielding 320 held-out models across the four methods. SpectralCanary achieves 94.8% average model-level accuracy,

showing that its per-generation discrimination signal supports dataset-use decisions at the model level. Appx. H provides the detailed protocol and TPR, FPR, precision, and F1 results.

4.3 The Same Setting Works Across Five Image Types

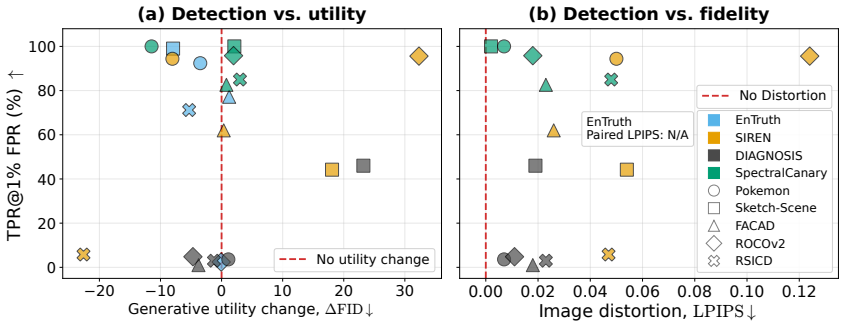

Figure 4: **SpectralCanary works consistently across five image types.** Each point shows one method on one dataset. The vertical axis is the per-generation TPR at 1% FPR: the percentage of outputs from protected-data models detected while 1% of outputs from clean-data models are falsely flagged. In (a), the horizontal axis measures the change in the trained model's output quality; zero means no change. In (b), it measures the change from an original image to its protected copy; zero means identical images. The best results are near the upper-left corner.

We next test whether the same SpectralCanary setting works beyond Pokemon by asking three questions: Does a single generated output support reliable per-generation discrimination? Do the protected images still resemble the originals? Does training on those images degrade model outputs?

SpectralCanary performs well on all three questions. At a per-generation FPR of 1%, it achieves an average per-generation TPR of 92.4% across outputs from models trained on protected data. The protected images remain close to their originals: average LPIPS is 0.022, where zero means identical images. Training on the protected images causes almost no average change in model output quality: mean ΔFID is 0.05, where zero means no change. On the sparse Sketch-Scene drawings, SpectralCanary achieves 100% per-generation TPR with LPIPS 0.002. The hardest dataset is RSICD, where it achieves 78.9% per-generation TPR with LPIPS 0.048.

Fig. 4 compares all methods. Each point shows one method on one dataset. Good results appear near the upper-left corner: high per-generation TPR with little change to either the training images or the trained model's outputs. SpectralCanary's green points remain near this region across all five datasets. DIAGNOSIS changes the original images little but usually provides weak per-generation discrimination. SIREN works well on some datasets but poorly on others and makes the largest image changes. EnTruth works well on several datasets but fails on ROCOv2. It has no LPIPS value because it adds new canary images rather than making a protected copy of each original. Fig. 1(c) and Appx. N show protected images beside their originals; Appx. B reports all numerical results.

4.4 Protection Is Fast as Well as Effective

SpectralCanary protects an image in 0.39 seconds on CPU while retaining the highest average per-generation TPR. Fig. 5 compares per-generation discrimination with generated-image quality, protection time, and image distortion on Pokemon. DIAGNOSIS is faster, but its average TPR is only 32.2%. SIREN and EnTruth require substantially more computation while achieving lower per-generation TPR. SpectralCanary also has no average FID penalty and the lowest paired-image LPIPS among the methods that modify every image. These timings are hardware-dependent, and CPU and GPU measurements are not perfectly interchangeable; their main purpose is to distinguish a single deterministic FFT-based transform from methods that require generating auxiliary images or performing iterative optimization. Appx. E gives the complete timing results.

4.5 Low Frequencies Carry the Trace More Reliably

The advantage of low frequencies remains after controlling for the amount of spectral change. We compare disjoint low-, mid-, and high-frequency bands on Pokemon and ROCOv2 with SD1.5.

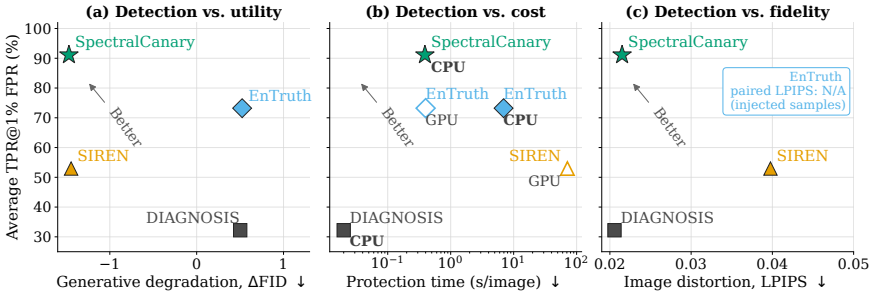

Figure 5: **SpectralCanary achieves high per-generation discrimination with low quality and runtime cost.** Average per-generation TPR at 1% FPR over the eight Pokemon settings is compared with (a) change in FID, (b) protection time, and (c) paired-image LPIPS. Preferred results lie upward and to the left. EnTruth adds 0.5% auxiliary samples, so paired-image LPIPS is not defined for it.

The bands contain the same number of Fourier coefficients, and we give them the same total amount of Fourier-domain change as the deployed smooth low-pass mask. This comparison prevents one band from winning merely because it changes more of the image. Fig. 6 shows that the low-frequency band has the highest per-generation TPR on both datasets. It exceeds the high-frequency band by 16.9 percentage points on Pokemon and 88.9 points on ROCOv2; the mid-frequency band is also weaker in both cases. Appx. K describes the construction, and Tab. 6 reports per-generation discrimination, image fidelity, generated-image quality, and variation across three canary images.

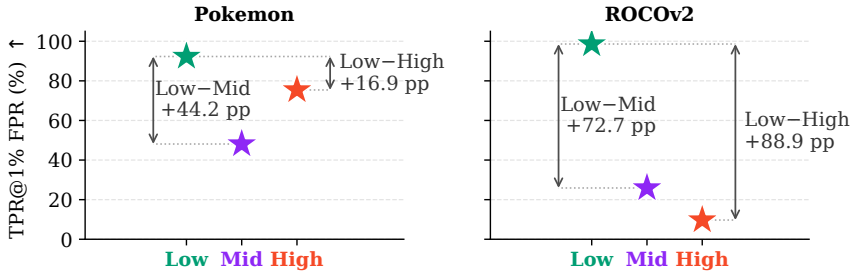

Figure 6: **Matched low-frequency traces transfer more reliably than mid- or high-frequency traces.** Points show mean per-generation TPR at 1% FPR over three canary images. Arrows show the performance gaps.

4.6 Removal Attacks Damage Quality Before Erasing the Trace

In our evaluated attacks, substantially reducing per-generation TPR also substantially damages the training data or the resulting model. The attacker knows the protection algorithm and may know α and c, but not the canary, original images, or verifier. We test three attacks. A *donor* is an unrelated, unprotected image whose low-frequency content is used as a replacement source. Fixed-donor replacement uses the same donor for all protected images. Per-image replacement uses a different donor for each image. The third attack estimates the shared low-frequency pattern directly from the protected collection and removes it. Appx. M gives the equations and settings.

Fig. 7 follows each attack from mild to aggressive removal. Mild attacks leave per-generation TPR largely intact. Driving TPR toward the clean-training baseline requires more than 25 additional FID points under the strongest evaluated attack and also produces large LPIPS changes. This establishes a removal–quality boundary for these three attacks; it is not a guarantee against every possible attack, especially one that can inspect the model or verifier.

5 Limitations and Responsible Use

Our evaluation covers two T2I architecture families, five datasets, and both LoRA and full-parameter fine-tuning. It does not establish transfer to every model, or mixture of a protected collection with a

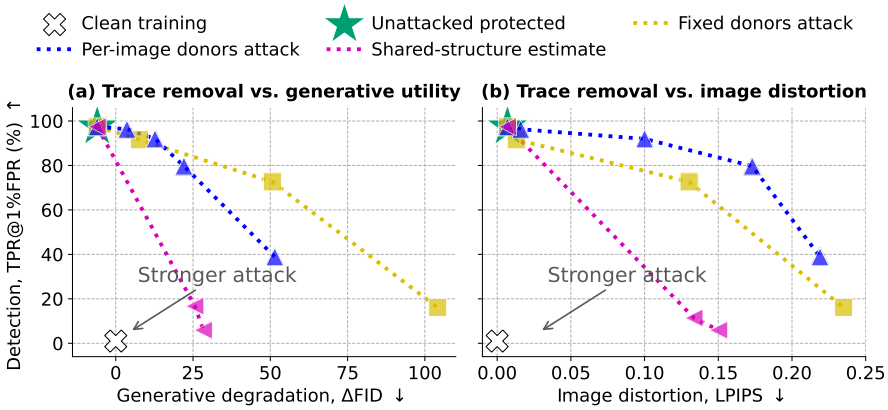

Figure 7: **The evaluated attacks erase the trace only after causing substantial damage.** From the unattacked point, stronger attacks move down and to the right, indicating per-generation TPR falls, but ΔFID and LPIPS increase. The clean-training marker shows the target TPR.

larger unprotected corpus. The main experiments also use one fixed natural canary image. Although the phase-randomized controls in Appx. L isolate the role of spatial coherence, broader evaluation across independently chosen canary images and key reuse is needed.

The adaptive evaluation assumes that the canary image and verifier remain private. If the attacker learns V, α, and c, Eq. 6 permits exact inversion before clipping or quantization. The three evaluated attacks (fixed-donor replacement, per-image donor replacement, and shared-structure estimation) characterize only this secrecy-dependent threat model and do not provide a cryptographic or white-box robustness guarantee. Future evaluations should consider collusion across releases, learned removal models, stronger shared-structure estimators, and adaptive access to verifier feedback.

SSIM, LPIPS, and FID capture only part of image fidelity and generative utility. Human perceptual studies and task-specific utility measurements would be needed to support stronger claims about invisibility or harmlessness. Likewise, our KS-test rejection rates on protected-data models estimate statistical power under the stated sampling protocol, while fixed-threshold model-level verification uses separate clean-data and protected-data models for calibration. A deployment would require broader calibration across representative clean-data models and should treat a positive result as statistical evidence under the threat model, not as standalone proof of unauthorized use.

SpectralCanary is intended to help data owners audit model customization, but proactive traces can also be inserted into data without the knowledge of downstream users. Responsible use requires authorization to modify the collection, disclosure where appropriate, respect for dataset licenses, and safeguards against unsupported attribution. Decisions with legal or reputational consequences should combine the audit with independent evidence.

6 Conclusion

We propose SpectralCanary, a deterministic and inexpensive approach to proactive T2I dataset-use auditing. SpectralCanary injects a weak shared visual trace through a smooth low-pass mask while leaving captions unchanged. Across eight Pokemon model–transformation settings, it provides the strongest average per-generation discrimination and reliable KS-based multi-query auditing among the evaluated methods. The same operating point generalizes across five visual domains while preserving protected-image fidelity and generative utility. Matched-energy frequency-location experiments further show that low-frequency placement transfers more consistently than mid/high-frequency placement across datasets. Finally, the adaptive evaluation identifies a scoped robustness boundary, showing that the evaluated sanitization attacks substantially reduce per-generation TPR only at the cost of considerable image-quality or generative-utility degradation. Together, these results establish SpectralCanary as a practical method for T2I dataset-use auditing, while the limitations discussed above motivate future work.

Reproducibility Statement

Sec. 2 specifies the algorithms. Appx. A reports the datasets, model and baseline settings, verifier training, data transformations, generation parameters, metrics, and hypothesis-testing protocol. Each main setting uses three independent training and generation runs; deterministic fidelity measurements are reported separately. The supplementary file also includes pytorch-like code for reproducibility purpose.

AI Use Disclosure

Generative AI assisted with language editing, manuscript organization, title development, presentation of existing results, and editing experimental code. It did not generate data, derive results, create citations, formulate hypotheses or mathematical claims, design the experimental methodology, or interpret experimental results. The authors verified all revisions and take responsibility for the text, claims, and artifacts.

References

Namhyuk Ahn, KiYoon Yoo, Wonhyuk Ahn, Daesik Kim, and Seung-Hun Nam. Nearly zero-cost protection against mimicry by personalized diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 28801–28810, 2025.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.

Yuetian Chen, Kaiyuan Zhang, Yuntao Du, Edoardo Stoppa, Charles Fleming, Ashish Kundu, Bruno Ribeiro, and Ninghui Li. Membership inference attacks against fine-tuned diffusion language models. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=oWKJursYpW.

Pinaki Nath Chowdhury, Aneeshan Sain, Ayan Kumar Bhunia, Tao Xiang, Yulia Gryaditskaya, and Yi-Zhe Song. Fs-coco: Towards understanding of freehand sketches of common objects in context. In *ECCV*, 2022.

Yingqian Cui, Jie Ren, Yuping Lin, Han Xu, Pengfei He, Yue Xing, Lingjuan Lyu, Wenqi Fan, Hui Liu, and Jiliang Tang. Ft-shield: A watermark against unauthorized fine-tuning in text-to-image diffusion models. *SIGKDD Explor. Newsl.*, 26(2):76–88, January 2025a. ISSN 1931-0145. doi: 10.1145/3715073.3715080. URL https://doi.org/10.1145/3715073.3715080.

Yingqian Cui, Jie Ren, Han Xu, Pengfei He, Hui Liu, Lichao Sun, Yue Xing, and Jiliang Tang. Diffusionshield: A watermark for data copyright protection against generative diffusion models. *SIGKDD Explor. Newsl.*, 26(2):60–75, January 2025b. ISSN 1931-0145. doi: 10.1145/3715073.3715079. URL https://doi.org/10.1145/3715073.3715079.

Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.

Jonas Geiping, Liam H Fowl, W. Ronny Huang, Wojciech Czaja, Gavin Taylor, Michael Moeller, and Tom Goldstein. Witches' brew: Industrial scale data poisoning via gradient matching. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=01olnfLIbD.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Sy-Tuyen Ho, Yaseema Rusiru Ariyarathna Epa, Yasoda Lasiru Ariyarathna Epa, Andrew Mendez, Kecheng Liu, Xudong Jiang, Alex Kot, Furong Huang, and Ngai-Man Cheung. Memory makes the poison: Over memorization drives visual poisoning in LVLMs, 2025. URL https://openreview.net/forum?id=2HGL1Szcp2.

Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.

Sangwon Jang, June Suk Choi, Jaehyeong Jo, Kimin Lee, and Sung Ju Hwang. Silent branding attack: Trigger-free data poisoning attack on text-to-image diffusion models. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8203–8212. IEEE, 2025.

Wojciech Łapacz, Stanisław Pawlak, Jan Dubiński, Franziska Boenisch, and Adam Dziedzic. Dataset usage inference without shadow models or held-out data. *arXiv preprint arXiv: 2606.26257*, 2026.

Boheng Li, Yanhao Wei, Yankai Fu, Zhenting Wang, Yiming Li, Jie Zhang, Run Wang, and Tianwei Zhang. Towards Reliable Verification of Unauthorized Data Usage in Personalized Text-to-Image Diffusion Models . In *2025 IEEE Symposium on Security and Privacy (SP)*, pp. 2564–2582, Los Alamitos, CA, USA, May 2025. IEEE Computer Society. doi: 10.1109/SP61157.2025.00073. URL https://doi.ieeecomputersociety.org/10.1109/SP61157.2025.00073.

Chumeng Liang and Jiaxuan You. Real-world benchmarks make membership inference attacks fail on diffusion models. *arXiv preprint arXiv:2410.03640*, 2024.

Yixin Liu, Chenrui Fan, Yutong Dai, Xun Chen, Pan Zhou, and Lichao Sun. Metacloak: Preventing unauthorized subject-driven text-to-image diffusion-based synthesis via meta-learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 24219–24228, June 2024.

Xiaoqiang Lu, Binqiang Wang, Xiangtao Zheng, and Xuelong Li. Exploring models and data for remote sensing image caption generation. *IEEE Transactions on Geoscience and Remote Sensing*, 56(4):2183–2195, 2017.

Zehong Ma, Longhui Wei, Shuai Wang, Shiliang Zhang, and Qi Tian. Deco: Frequency-decoupled pixel diffusion for end-to-end image generation. In *Computer Vision and Pattern Recognition (CVPR)*, 2026. URL https://zehong-ma.github.io/DeCo/.

Tuan Anh Nguyen and Anh Tuan Tran. Wanet - imperceptible warping-based backdoor attack. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=eEn8KTtJOx.

Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=a68SUt6zFt. Featured Certification.

Zhuoshi Pan, Yuguang Yao, Gaowen Liu, Bingquan Shen, H. Vicky Zhao, Ramana Rao Kompella, and Sijia Liu. From trojan horses to castle walls: Unveiling bilateral data poisoning effects in diffusion models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=yiXZZC5qDI.

Yan Pang and Tianhao Wang. Black-box membership inference attacks against fine-tuned diffusion models. In *NDSS*, 2025.

Justin N. M. Pinkney. Pokemon blip captions. https://huggingface.co/datasets/lambdalabs/pokemon-blip-captions/, 2022.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PmLR, 2021.

Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International conference on machine learning*, pp. 5301–5310. PMLR, 2019.

Jie Ren, Yingqian Cui, Chen Chen, Yue Xing, Hui Liu, and Lingjuan Lyu. Entruth: Tracing the unauthorized dataset usage in diffusion models. In *Generative and Protective AI for Content Creation*, 2025. URL https://openreview.net/forum?id=mbreZbKz6m.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, June 2022.

Johannes Rückert, Louise Bloch, Raphael Brüngel, Ahmad Idrissi-Yaghir, Henning Schäfer, Cynthia S Schmidt, Sven Koitka, Obioma Pelka, Asma Ben Abacha, Alba G. Seco de Herrera, et al. Rocov2: Radiology objects in context version 2, an updated multimodal image dataset. *Scientific Data*, 11(1):688, 2024.

Maximilian Seitzer. pytorch-fid: FID Score for PyTorch. https://github.com/mseitzer/pytorch-fid, August 2020. Version 0.3.0.

Ali Shafahi, W Ronny Huang, Mahyar Najibi, Octavian Suciu, Christoph Studer, Tudor Dumitras, and Tom Goldstein. Poison frogs! targeted clean-label poisoning attacks on neural networks. *Advances in neural information processing systems*, 31, 2018.

Shawn Shan, Jenna Cryan, Emily Wenger, Haitao Zheng, Rana Hanocka, and Ben Y. Zhao. Glaze: Protecting artists from style mimicry by Text-to-Image models. In *32nd USENIX Security Symposium (USENIX Security 23)*, pp. 2187–2204, Anaheim, CA, August 2023. USENIX Association. ISBN 978-1-939133-37-3. URL https://www.usenix.org/conference/usenixsecurity23/presentation/shan.

Shawn Shan, Wenxin Ding, Josephine Passananti, Stanley Wu, Haitao Zheng, and Ben Y Zhao. Nightshade: Prompt-specific poisoning attacks on text-to-image generative models. In *2024 IEEE symposium on security and privacy (SP)*, pp. 807–825. IEEE, 2024.

Thanh Van Le, Hao Phung, Thuan Hoang Nguyen, Quan Dao, Ngoc N Tran, and Anh Tran. Anti-dreambooth: Protecting users from personalized text-to-image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2116–2127, 2023.

Haonan Wang, Qianli Shen, Yao Tong, Yang Zhang, and Kenji Kawaguchi. The stronger the diffusion model, the easier the backdoor: Data poisoning to induce copyright breaches without adjusting finetuning pipeline. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 51465–51483. PMLR, 21–27 Jul 2024a. URL https://proceedings.mlr.press/v235/wang24bm.html.

Zhenting Wang, Chen Chen, Lingjuan Lyu, Dimitris N. Metaxas, and Shiqing Ma. DIAGNOSIS: Detecting unauthorized data usages in text-to-image diffusion models. In *The Twelfth International Conference on Learning Representations*, 2024b. URL https://openreview.net/forum?id=f8S3aLm0Vp.

Yanhao Wei, Xiaokang Zhao, Boheng Li, Yang Zhang, and Run Wang. RECOVER: Reliable detection of unauthorized data usage in text-to-image diffusion models via inversion robustness. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=HYZ3ooGuew.

Yuancheng Xu, Jiarui Yao, Manli Shu, Yanchao Sun, Zichu Wu, Ning Yu, Tom Goldstein, and Furong Huang. Shadowcast: Stealthy data poisoning attacks against vision-language models. *Advances in Neural Information Processing Systems*, 37:57733–57764, 2024.

Zhi-Qin John Xu, Yaoyu Zhang, Tao Luo, Yanyang Xiao, and Zheng Ma. Frequency principle: Fourier analysis sheds light on deep neural networks. *arXiv preprint arXiv:1901.06523*, 2019.

Xuewen Yang, Heming Zhang, Di Jin, Yingru Liu, Chi-Hao Wu, Jianchao Tan, Dongliang Xie, Jue Wang, and Xin Wang. Fashion captioning: Towards generating accurate descriptions with semantic rewards. In *European conference on computer vision*, pp. 1–17. Springer, 2020.

Mingyang Yi, Aoxue Li, Yi Xin, and Zhenguo Li. Towards understanding the working mechanism of text-to-image diffusion model. *Advances in Neural Information Processing Systems*, 37:55342–55369, 2024.

Shengfang Zhai, Yinpeng Dong, Qingni Shen, Shi Pu, Yuejian Fang, and Hang Su. Text-to-image diffusion models can be easily backdoored through multimodal data poisoning. In *Proceedings of the 31st ACM International Conference on Multimedia*, pp. 1577–1587, 2023.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.

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 *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024a. URL https://openreview.net/forum?id=7hy5fy2OC6.

Zhengyue Zhao, Jinhao Duan, Xing Hu, Kaidi Xu, Chenan Wang, Rui Zhang, Zidong Du, Qi Guo, and Yunji Chen. Unlearnable examples for diffusion models: Protect data from unauthorized exploitation. In *ICLR 2024 Workshop on Reliable and Responsible Foundation Models*, 2024b. URL https://openreview.net/forum?id=CHTRtLPYNF.

Jie Zhu and Leye Wang. Auditing data provenance in real-world text-to-image diffusion models for privacy and copyright protection. *arXiv preprint arXiv:2506.11434*, 2025.

APPENDIX

APPENDIX ROADMAP

The appendix is organized into five blocks:

- **Protocol and primary evidence:** Appx. A specifies the metrics and implementation, Appx. C and B give the complete benchmark tables, Appx. D reports hyperparameter sweeps, and Appx. E reports protection cost.
- **Additional training and auditing evidence:** Appx. F studies learning dynamics, Appx. G evaluates full-parameter fine-tuning, Appx. H reports fixed-threshold model-level verification, Appx. I shows generated samples, and Appx. J provides KS-based multi-query auditing results.
- **Mechanism and robustness analyses:** Appx. K and L study spectral location and canary structure, while M and N provide adaptive-attack details, and additional protected-image examples.
- **Baseline analyses:** Appx. O and P analyze EnTruth, while Appx. Q adapts Silent Branding.
- **Foundations and context:** Appx. R provides the theoretical analysis, and Appx. S gives extended related work.

A DETAILED EXPERIMENTAL PROTOCOL

A.1 AUDITING METRICS

Per-generation AUC and TPR at a target FPR. In each experiment setting, we train one clean-data LoRA $\mathcal{G}^-$ and one protected-data LoRA $\mathcal{G}^+$. We query both models using the same prompt list, generate one image per prompt, and score every generated image independently with the verifier. This produces the score sets $\mathcal{S}^- = \{s_j^{\mathcal{G}^-}\}_{j=1}^K$ and $\mathcal{S}^+ = \{s_j^{\mathcal{G}^+}\}_{j=1}^K$. The AUC measures separation between these distributions and equals the probability that an independently sampled generation

from the protected-data model receives a higher score than an independently sampled generation from the clean-data model, with ties counted as half correct:

$$\mathrm{AUC} = \Pr_{s^+\sim\mathcal{S}^+,\, s^-\sim\mathcal{S}^-}\left[s^+ > s^-\right] + \frac{1}{2}\Pr_{s^+\sim\mathcal{S}^+,\, s^-\sim\mathcal{S}^-}\left[s^+ = s^-\right]. \tag{7}$$

For a score threshold t, the empirical false-positive and true-positive rates are

$$\mathrm{FPR}(t) = \frac{1}{|\mathcal{S}^-|}\sum_{s\in\mathcal{S}^-}\mathbb{I}[s \geq t], \qquad \mathrm{TPR}(t) = \frac{1}{|\mathcal{S}^+|}\sum_{s\in\mathcal{S}^+}\mathbb{I}[s \geq t]. \tag{8}$$

At a target per-generation FPR r, we enumerate the empirical operating points induced by the observed score thresholds and select the feasible operating point with the largest FPR satisfying $\mathrm{FPR} \leq r$; ties are broken by larger TPR. We report the corresponding TPR at FPR targets of 5% and 1%.

KS-based multi-query auditing. Following Li et al. (2025), we evaluate whether suspect-model generations are statistically distinguishable from clean-data-model generations. For each setting, we uniformly sample one fixed clean reference set $\mathcal{R}$ of 300 clean-data-model scores. Each trial independently samples a suspect query set $\mathcal{Q}_b$ of $N_q = 30$ protected-data-model scores without replacement within that trial; scores may reappear across trials. The one-sided alternative follows the verifier's score direction. For methods in which lower scores indicate protected-data use, we use the alternative below; for methods in which higher scores indicate protected-data use, its direction is reversed.

$$H_0 : F_{\mathcal{Q}_b} = F_{\mathcal{R}}, \qquad H_1 : F_{\mathcal{Q}_b} > F_{\mathcal{R}}, \tag{9}$$

with statistic

$$D_b^+ = \sup_z\left[\widehat{F}_{\mathcal{Q}_b}(z) - \widehat{F}_{\mathcal{R}}(z)\right]. \tag{10}$$

For significance level γ, the KS rejection rate on protected-data models is

$$\mathrm{RejectRate}(\gamma) = \frac{1}{B}\sum_{b=1}^{B}\mathbf{1}[p_b < \gamma], \tag{11}$$

where $B = 100{,}000$. We perform the resampling independently for each of three runs and report the mean and standard deviation.

A.2 Implementation Details

Fine-tuning and auditing across methods. We follow Li et al. (2025) for the main setting and use at most $1{,}000$ training images. For each dataset, backbone, method, and transformation, we train one clean-data and one protected-data LoRA. SD1.5 uses resolution 512×512, batch size 4, LoRA rank 64, AdamW, learning rate 10^{-4}, and 80 epochs. Generation uses Euler ancestral sampling with 30 steps and classifier-free guidance 7.5. SD3.5 uses LoRA rank 16, learning rate 10^{-4}, AdamW, a constant schedule, and 80 epochs. Its objective uses the unpreconditioned linear flow-matching target $u = \epsilon - x_0$. SD3.5 generations use resolution 1024×1024 and 30 steps.

The clean-data and protected-data models receive the same ordered prompt list, and every generation is scored independently. SIREN and SpectralCanary use original captions as audit prompts; DIAGNOSIS appends its trigger token; EnTruth uses method-specific trigger prompts. SD1.5 experiments generate $1{,}000$ images per model and SD3.5 experiments generate 500.

EnTruth (Ren et al., 2025). As there is no official code, we follow its shared-template pipeline. Candidate templates use the prompts "billboard for big sale", "a painting with a frame", "photo frame with a family", "a window with mountains outside". Foregrounds use "fruits for sale" and "fruit stand", with 50 images per foreground prompt and one per template prompt, generated by SD1.5. The hard trigger is "[Tgr]", the soft trigger keyword is "fruit". The injected samples rate is 0.5%. An ImageNet-pretrained ResNet-50 verifier uses selected templated samples as positives and original samples as negatives. The main paper reports the stronger hard-trigger result from Appx. P.

DIAGNOSIS (Wang et al., 2024b). We follow the official code. DIAGNOSIS applies a WaNet-style image transformation (Nguyen & Tran, 2021) to 20% of the training set, uses hard trigger token "`tq`", and sets $k = 128$ and $s = 1.0$. An ImageNet-pretrained ResNet-50 classifier uses transformed images as positives and their originals as negatives.

SIREN (Li et al., 2025). We follow the official code. SIREN first trains a clean-data LoRA initialization, then fine-tunes its encoder–decoder for 60 epochs. The learned encoder–decoder protects the full dataset and directly scores generated samples. No trigger prompt is required.

SpectralCanary protection and export. We use the fixed RGB canary in Fig. 1(b). Protection is applied at the stored image resolution before downstream resizing or transformations. We convert to sRGB, scale channels to $[0, 1]$, resize the canary with antialiased bicubic interpolation, apply the unitary channel-wise transform, discard numerical imaginary residue, clip to $[0, 1]$, quantize to 8-bit RGB, and save lossless PNG files. Defaults are $\delta = 1$, $\alpha = 0.02$, and $c = 0.20$. For $\delta < 1$, the run seed selects subset $\mathcal{I}_\delta$ for both protection and verifier training.

SpectralCanary verifier. We initialize an ImageNet-pretrained ResNet-50 (He et al., 2016) with a binary head. Source indices are split $80/20$ before pairing, so each original and protected counterpart stays in the same partition. Inputs are resized to 224×224 and ImageNet-normalized without stochastic augmentation. We fine-tune all layers with AdamW, learning rate 10^{-4}, weight decay 10^{-4}, batch size 32, and 20 epochs, selecting the lowest validation-loss checkpoint. Each dataset and run trains a separate verifier without T2I generations.

Data transformations.

- **Training-time augmentation.** Images first pass through the bucketed training resize and are randomly cropped when larger than the target resolution. With probability 0.33, color augmentation applies either a hue shift in $[-8, 8]$ degrees or gamma adjustment with $\gamma \sim U(0.95, 1.05)$. Horizontal flipping has probability 0.5.
- **Prompt recaptioning.** Every caption is rewritten using Qwen2.5-VL 7B Instruct; images remain unchanged. The prompt is:

```
Rewrite the caption for this image as a high-quality caption.
Keep the important visual content, remove redundant wording,
and avoid mentioning text quality, watermarks, camera metadata,
or that this is a rewrite.  Output only the rewritten caption.
Current caption:  {current_caption}
```

- **Image regeneration.** WatermarkAttacker uses the `sd2-community` Stable Diffusion 2.1 checkpoint, VAE encoding, `noise_step=60`, 50 denoising steps, guidance 7.5, an empty prompt, and fixed seed 1024. It preserves the source resolution. The regenerated datasets train suspect LoRAs under the same downstream protocol.

B Full Cross-Domain Results

Fig. 4 gives the summary results. Tab. 1 reports the corresponding numerical results, including per-generation AUC, TPR at 5% FPR, KS-test rejection rates on protected-data models, generative utility, and image fidelity.

C Full Pokemon Benchmark Results

Tab. 2 expands the per-generation discrimination summary in Fig. 2 into its architecture-specific AUC and TPR results at 5% and 1% FPR.

D Protection-Hyperparameter Ablations

We vary blending strength α, protected fraction δ, and Gaussian bandwidth c on Pokemon with SD1.5. When $\delta < 1$, we use a random subset and paired-verifier protocol in Alg. 1. Fig. 8 shows

Table 1: SD3.5 comparison across datasets. Paired training-image fidelity is deterministic for a fixed protected dataset. EnTruth injects 0.5% auxiliary samples and therefore has no directly comparable paired-LPIPS value.

Dataset	Method	Per-Generation Discrimination (%) ↑			KS Rejection Rate on Protected-Data Models (%) ↑			Utility ΔFID ↓	Fidelity	
		AUC	TPR@5% FPR	TPR@1% FPR	$\gamma = 10^{-5}$	$\gamma = 10^{-9}$	$\gamma = 10^{-15}$		SSIM ↑	LPIPS ↓
Pokemon	EnTruth	96.5 ± 1.1	94.9 ± 1.7	93.9 ± 1.1	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	-3.06 ± 0.60	N/A	N/A
	SIREN	99.1 ± 0.1	96.9 ± 0.7	91.9 ± 1.8	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	-6.58 ± 1.04	0.980	0.050
	DIAGNOSIS	53.6 ± 1.5	8.2 ± 0.4	2.3 ± 1.1	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	3.65 ± 2.54	0.994	0.007
	SpectralCanary	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	-9.00 ± 1.82	0.998	0.007
Sketch-Scene	EnTruth	99.8 ± 0.1	98.8 ± 0.7	97.5 ± 1.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	-5.98 ± 1.50	N/A	N/A
	SIREN	97.5 ± 0.6	90.6 ± 1.3	44.2 ± 21.3	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	19.06 ± 0.76	0.975	0.054
	DIAGNOSIS	96.0 ± 1.4	80.3 ± 7.9	56.5 ± 12.0	100.0 ± 0.0	100.0 ± 0.0	87.7 ± 11.2	21.75 ± 1.04	0.980	0.019
	SpectralCanary	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	4.76 ± 2.53	0.997	0.002
FACAD	EnTruth	98.0 ± 0.7	91.1 ± 4.1	74.3 ± 6.5	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	1.15 ± 0.32	N/A	N/A
	SIREN	96.2 ± 0.3	80.8 ± 2.0	57.8 ± 6.2	100.0 ± 0.0	100.0 ± 0.0	99.0 ± 0.1	1.42 ± 0.76	0.985	0.026
	DIAGNOSIS	52.8 ± 2.0	5.9 ± 0.3	1.3 ± 0.2	0.0 ± 0.1	0.0 ± 0.0	0.0 ± 0.0	-3.02 ± 1.08	0.992	0.018
	SpectralCanary	99.4 ± 0.1	99.7 ± 0.4	86.0 ± 2.4	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	1.17 ± 0.40	0.986	0.023
ROCOv2	EnTruth	59.3 ± 1.7	12.5 ± 2.9	2.7 ± 0.8	0.4 ± 0.1	0.0 ± 0.0	0.0 ± 0.0	2.51 ± 1.82	N/A	N/A
	SIREN	99.9 ± 0.0	100.0 ± 0.0	96.8 ± 1.6	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	22.30 ± 7.08	0.909	0.124
	DIAGNOSIS	63.4 ± 1.1	15.0 ± 1.1	4.8 ± 0.8	2.1 ± 0.5	0.0 ± 0.0	0.0 ± 0.0	-5.27 ± 0.46	0.992	0.011
	SpectralCanary	99.8 ± 0.1	99.3 ± 0.3	97.1 ± 1.6	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	1.90 ± 1.06	0.976	0.028
RSICD	EnTruth	92.2 ± 5.6	68.7 ± 21.0	44.7 ± 24.0	98.3 ± 2.4	76.6 ± 33.1	66.0 ± 46.5	-18.00 ± 9.06	N/A	N/A
	SIREN	71.7 ± 1.4	10.5 ± 0.5	3.9 ± 1.9	41.3 ± 11.3	0.0 ± 0.1	0.0 ± 0.0	-21.10 ± 1.17	0.975	0.047
	DIAGNOSIS	55.6 ± 1.3	11.2 ± 1.8	4.5 ± 2.1	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	-2.70 ± 1.12	0.975	0.023
	SpectralCanary	98.1 ± 0.8	92.3 ± 2.8	78.9 ± 4.4	100.0 ± 0.0	100.0 ± 0.0	99.7 ± 0.4	1.44 ± 1.33	0.986	0.048
Average	EnTruth	89.2	73.2	62.6	79.7	75.3	73.2	-4.68	N/A	N/A
	SIREN	92.9	75.8	58.9	88.3	80.0	79.8	3.02	0.965	0.060
	DIAGNOSIS	64.3	24.1	13.9	20.4	20.0	17.5	2.88	0.987	0.016
	SpectralCanary	99.5	98.3	92.4	100.0	100.0	99.9	0.05	0.989	0.022

that the canary becomes increasingly visible as α grows. Fig. 9 shows improved discrimination with stronger blending and a larger protected fraction, while the bandwidth sweep saturates near the default $c = 0.20$. These sweeps motivate the default $\alpha = 0.02$, $\delta = 1$, and $c = 0.20$ operating point.

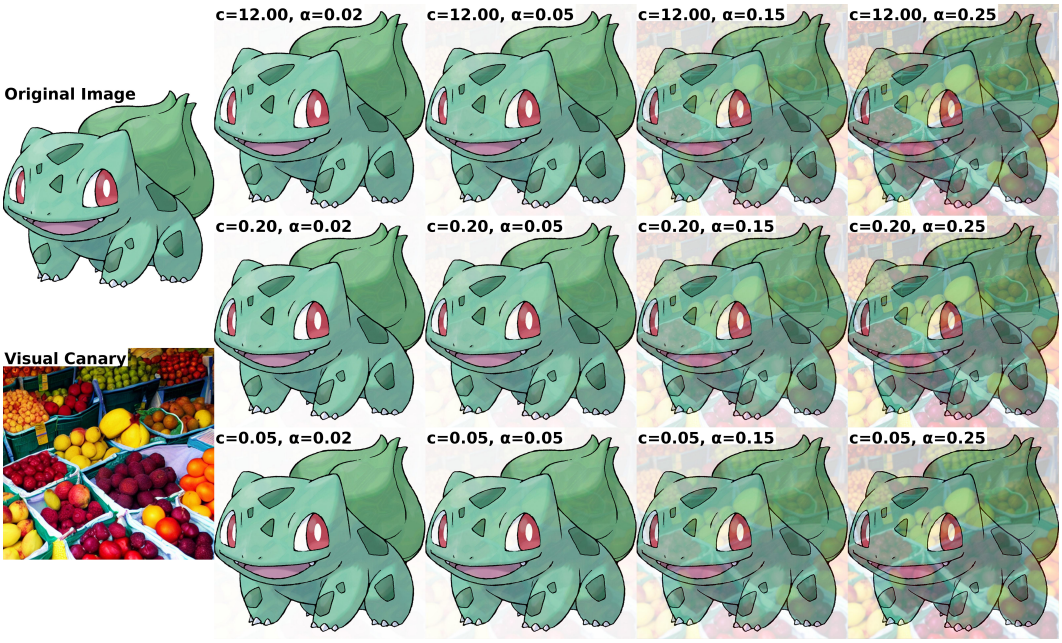

Figure 8: Protected images under different blending strengths α and Gaussian bandwidths c.

Table 2: **Per-generation discrimination on Pokemon.** Means and standard deviations over three runs.

Transformation	Model	Method	AUC (%) ↑	TPR@5% FPR (%) ↑	TPR@1% FPR (%) ↑
None	SD1.5	EnTruth	98.8 ± 0.2	94.1 ± 0.5	85.3 ± 2.6
		SIREN	98.4 ± 0.1	93.9 ± 0.3	66.3 ± 1.1
		DIAGNOSIS	99.7 ± 0.2	100 ± 0.0	99.9 ± 0.1
		SpectralCanary	99.7 ± 0.1	98.8 ± 0.4	97.7 ± 0.5
	SD3.5	EnTruth	96.5 ± 1.1	94.9 ± 1.7	93.9 ± 1.1
		SIREN	99.1 ± 0.1	96.9 ± 0.7	91.9 ± 1.8
		DIAGNOSIS	53.6 ± 1.5	8.2 ± 0.4	2.3 ± 1.1
		SpectralCanary	100 ± 0.0	100 ± 0.0	100 ± 0.0
Training-Time Augmentation	SD1.5	EnTruth	97.4 ± 0.9	91.8 ± 2.2	76.2 ± 6.3
		SIREN	64.0 ± 0.8	18.1 ± 1.3	6.9 ± 1.4
		DIAGNOSIS	99.6 ± 0.4	100 ± 0.0	88.3 ± 10.4
		SpectralCanary	94.9 ± 1.5	85.5 ± 7.8	79.2 ± 12.3
	SD3.5	EnTruth	99.1 ± 0.6	97.3 ± 1.7	95.3 ± 1.6
		SIREN	80.0 ± 3.2	40.2 ± 2.1	18.8 ± 1.7
		DIAGNOSIS	76.0 ± 0.5	25.3 ± 0.3	8.2 ± 2.6
		SpectralCanary	91.3 ± 8.9	76.6 ± 20.4	65.9 ± 22.4
Prompt Recaptioning	SD1.5	EnTruth	75.7 ± 1.1	34.9 ± 1.5	19.4 ± 2.2
		SIREN	97.0 ± 0.8	82.5 ± 1.1	63.6 ± 1.0
		DIAGNOSIS	35.3 ± 1.8	2.4 ± 0.2	0.4 ± 0.3
		SpectralCanary	99.2 ± 0.2	96.0 ± 0.4	93.6 ± 1.6
	SD3.5	EnTruth	83.6 ± 1.7	52.3 ± 3.1	35.9 ± 1.6
		SIREN	99.0 ± 0.2	97.6 ± 0.9	92.0 ± 1.1
		DIAGNOSIS	75.0 ± 0.9	30.4 ± 5.8	13.0 ± 7.9
		SpectralCanary	100 ± 0.0	100 ± 0.0	100 ± 0.0
Image Regeneration	SD1.5	EnTruth	99.1 ± 0.3	95.7 ± 1.0	88.1 ± 3.2
		SIREN	93.5 ± 1.0	70.0 ± 0.9	44.4 ± 2.4
		DIAGNOSIS	98.8 ± 0.5	100 ± 0.0	43.9 ± 24.6
		SpectralCanary	99.2 ± 0.3	95.3 ± 1.0	92.5 ± 1.2
	SD3.5	EnTruth	95.4 ± 2.4	93.1 ± 3.1	91.7 ± 3.4
		SIREN	85.0 ± 1.3	58.0 ± 2.1	37.2 ± 2.1
		DIAGNOSIS	59.6 ± 1.5	7.5 ± 0.7	1.7 ± 0.4
		SpectralCanary	100 ± 0.0	100 ± 0.0	100 ± 0.0
Average over all eight settings		EnTruth	93.2	81.8	73.2
		SIREN	89.5	69.7	52.6
		DIAGNOSIS	74.7	46.7	32.2
		SpectralCanary	**98.0**	**94.0**	**91.1**

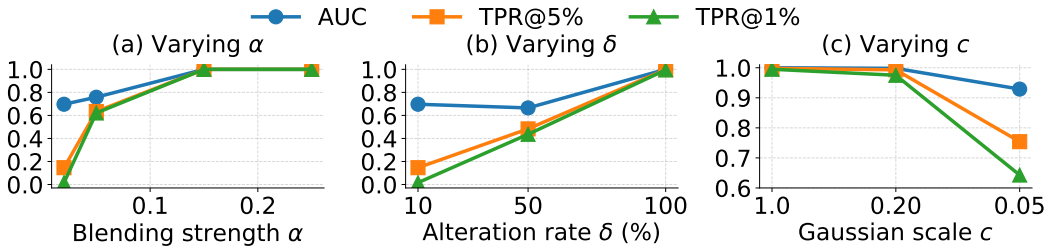

Figure 9: Protection-hyperparameter ablations on Pokemon with SD1.5: (a) blending strength, fixing $c = 0.2, \delta = 10\%$; (b) protected fraction, fixing $c = 0.2, \alpha = 0.02$; and (c) Gaussian bandwidth, fixing $\delta = 100\%, \alpha = 0.02$

Additional Evidence and Robustness Analyses

E Dataset Protection Cost

SpectralCanary protects an image with one FFT-based blend followed by an inverse transform. It requires neither iterative optimization nor an auxiliary generative model. Because the baselines use

Table 3: End-to-end dataset protection cost and downstream per-generation discrimination on Pokemon. Protection time includes every stage required to produce the protected dataset and is amortized over samples. We omit CPU measurements for SIREN and Silent Branding because their pipelines require GPU-based optimization. DIAGNOSIS and SpectralCanary require no GPU for protection and therefore do not benefit from GPU acceleration. KS-based multi-query auditing uses 30 generated images per test.

	Protection time (s/img) ↓		Per-Generation Discrimination (%) ↑			KS-Based Multi-Query Audit Rejection Rate (%) ↑		
Method	CPU	GPU	AUC	TPR@5% FPR	TPR@1% FPR	$\gamma = 10^{-5}$	$\gamma = 10^{-9}$	$\gamma = 10^{-15}$
DIAGNOSIS	**0.02**	**0.02**	74.7	46.7	32.2	49.7	38.2	37.5
SIREN	–	71.38	89.5	69.7	52.6	84.8	70.8	57.2
EnTruth	6.92	0.40	93.2	81.8	73.2	93.3	80.8	75.2
Silent Branding	–	273.13	80.9	32.5	12.8	73.8	34.1	2.7
SpectralCanary	0.39	0.39	**98.0**	**94.0**	**91.1**	**99.6**	**96.4**	**83.6**

different protection pipelines, operation counts are not directly comparable, where SIREN optimizes a learned watermark, Silent Branding fine-tunes a DreamBooth model and performs region selection and iterative refinement, and EnTruth generates and filters auxiliary samples. We therefore measure the end-to-end time required to produce each protected dataset on identical hardware. GPU measurements use one NVIDIA RTX A6000 with eight CPU cores; CPU-only measurements use the same system with the GPU disabled.

Tab. 3 reports 0.39 seconds per image for SpectralCanary on CPU. The remaining columns repeat the corresponding averages from the main Pokemon evaluation so that protection cost can be compared with downstream per-generation discrimination.

F The Trace Appears Early in Fine-Tuning

The main experiments follow Li et al. (2025) and fine-tune each method for 80 epochs. To measure when the trace appears, we apply the verifier to generations from intermediate checkpoints and report per-generation AUC and TPR at a target FPR. At each checkpoint, we generate 100 matched images from the clean-data and protected-data models using the same captions. In Fig. 10, the base models at epoch 0 are near chance (AUC ≈ 0.500). After five epochs, AUC reaches 0.956 and TPR at 1% FPR reaches 0.730; performance approaches saturation by epochs 10–15 and remains similar through epoch 80.

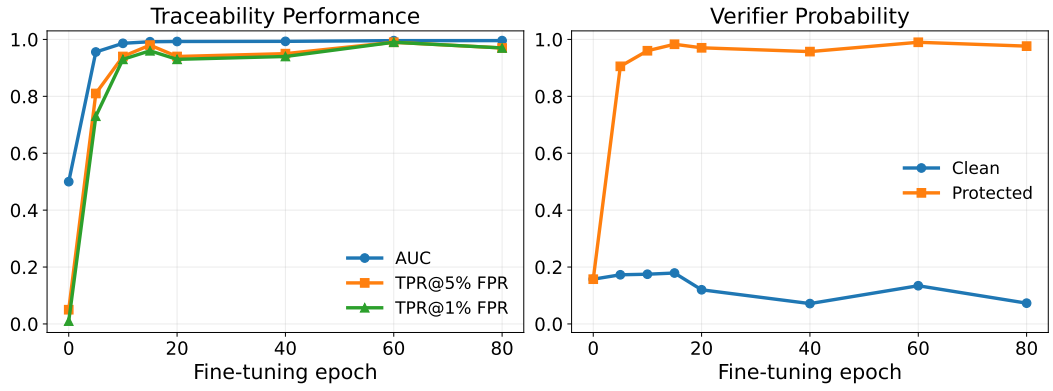

Figure 10: **Per-generation discrimination emerges early and remains stable during fine-tuning.** Using a defender's verifier, we report (left) per-generation AUC and TPR at FPRs of 5% and 1%, and (right) the mean predicted protected-class probability for clean-data-model and protected-data-model generations. Each checkpoint uses 100 matched Pokemon generations per condition with the same captions.

G Full Fine-Tuning Evaluation

Table 4: **SpectralCanary under full SD1.5 fine-tuning on Pokemon.** We report mean and standard deviation across three runs.

Transformation	Per-Generation Discrimination (%) ↑			KS Rejection Rate on Protected-Data Models (%) ↑		
	AUC	TPR@5% FPR	TPR@1% FPR	$\gamma = 10^{-5}$	$\gamma = 10^{-9}$	$\gamma = 10^{-15}$
None	98.8 ± 0.3	94.0 ± 1.7	91.7 ± 2.9	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
Training-Time Augmentation	84.3 ± 4.9	66.4 ± 10.2	61.9 ± 12.0	91.7 ± 11.6	68.2 ± 36.3	30.0 ± 27.9
Prompt Recaptioning	97.9 ± 0.2	88.8 ± 1.5	85.5 ± 3.0	100.0 ± 0.0	100.0 ± 0.0	99.5 ± 0.2
Image Regeneration	99.7 ± 0.1	97.8 ± 0.8	96.0 ± 1.6	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
Average	95.2	86.7	83.8	97.9	92.1	82.3

The main paper evaluates SpectralCanary with LoRA fine-tuning, which is the most practical customization setting for T2I models because it requires substantially less memory and storage than updating all base-model parameters. To test robustness of the trace beyond LoRA, we additionally evaluate **full-parameter fine-tuning** of SD1.5 on Pokemon. We follow the same SD1.5/Pokemon setup but train the T2I models for 50 epochs.

Tab. 4 shows that SpectralCanary remains effective under full-parameter fine-tuning. It achieves 98.8% AUC, 91.7% TPR at 1% FPR, and an essentially perfect KS rejection rate on protected-data models even at $\gamma = 10^{-15}$. Prompt recaptioning and image regeneration remain similarly robust. Training-time augmentation is the most challenging setting, reducing TPR at 1% FPR to 61.9% and the KS rejection rate at $\gamma = 10^{-15}$ to 30.0%, but the trace remains detectable at less stringent levels. Across all four settings, the average AUC is 95.2% and the average TPR at 1% FPR is 83.8%.

H Model-Level Verification

Our per-generation discrimination evaluation measures how well verifier scores separate generations from clean-data and protected-data models. This protocol enables a controlled comparison across methods without threshold calibration or query aggregation. Fixed-threshold model-level verification instead translates aggregated per-generation scores into a model-level decision that answers: *was a suspect model trained on the protected dataset?*

We first calibrate a model-level decision threshold and then keep it fixed for all held-out evaluations. For each method and T2I architecture, we reserve one independently trained clean-data model $\mathcal{G}^-_{\text{cal}}$ and one protected-data model $\mathcal{G}^+_{\text{cal}}$ for threshold calibration.

Given K audit prompts, we generate one image per prompt and define the aggregate score of a model $\mathcal{G}$ as

$$q_K(\mathcal{G}) = \frac{1}{K} \sum_{j=1}^{K} f_\theta\big(y_j^{\mathcal{G}}\big)\,, \tag{12}$$

where f_θ is the method-specific verifier and $y_j^{\mathcal{G}}$ is the image generated for prompt j.

For calibration, we use all available generations with $K_{\text{cal}} = 1000$ for SD1.5 and $K_{\text{cal}} = 500$ for SD3.5. We set the method–architecture-specific threshold to the midpoint between the clean-data and protected-data calibration scores:

$$t = \frac{q_{K_{\text{cal}}}(\mathcal{G}^-_{\text{cal}}) + q_{K_{\text{cal}}}(\mathcal{G}^+_{\text{cal}})}{2}. \tag{13}$$

The midpoint provides a simple symmetric decision boundary between the two calibration models without favoring either class. The calibration models are excluded from evaluation. Each threshold is calibrated once under the no-transformation setting and then held fixed across all transformation settings, so evaluation does not assume prior knowledge of the transformation applied by the unauthorized trainer.

Table 5: **Thirty-query fixed-threshold model-level verification across transformation settings.** Thresholds are calibrated once from all 1,000 SD1.5 or 500 SD3.5 queries of separate no-transformation calibration models. Each suspect model is evaluated using 30 queries. We repeat the 30-query sampling 1,000 times from the held-out generations and report the mean and standard deviation of model-level metrics across these resampled query sets.

Transformations	Model	Method	TPR (%) ↑	FPR (%) ↓	Precision (%) ↑	F1 (%) ↑	Accuracy (%) ↑
None	SD1.5	EnTruth	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
		SIREN	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
		DIAGNOSIS	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
		SpectralCanary	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
	SD3.5	EnTruth	64.1 ± 8.1	0.0 ± 0.0	100.0 ± 0.0	77.8 ± 5.6	82.0 ± 4.0
		SIREN	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
		DIAGNOSIS	55.0 ± 17.0	27.5 ± 14.9	67.5 ± 16.1	59.4 ± 14.6	63.7 ± 11.5
		SpectralCanary	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
Training-Time Augmentation	SD1.5	EnTruth	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
		SIREN	4.7 ± 9.3	0.0 ± 0.0	21.6 ± 41.2	7.6 ± 14.9	52.3 ± 4.6
		DIAGNOSIS	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
		SpectralCanary	99.2 ± 3.8	0.0 ± 0.0	100.0 ± 0.0	99.6 ± 2.1	99.6 ± 1.9
	SD3.5	EnTruth	80.7 ± 3.7	0.0 ± 0.0	100.0 ± 0.0	89.3 ± 2.0	90.3 ± 1.8
		SIREN	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	50.0 ± 0.0
		DIAGNOSIS	94.6 ± 9.8	17.4 ± 12.8	85.4 ± 10.1	89.3 ± 7.9	88.6 ± 8.3
		SpectralCanary	17.0 ± 10.5	0.0 ± 0.0	77.5 ± 41.8	27.6 ± 16.1	58.5 ± 5.3
Prompt Recaptioning	SD1.5	EnTruth	57.4 ± 13.7	0.0 ± 0.0	100.0 ± 0.0	72.0 ± 11.1	78.7 ± 6.9
		SIREN	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
		DIAGNOSIS	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	50.0 ± 0.0
		SpectralCanary	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
	SD3.5	EnTruth	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	50.0 ± 0.0
		SIREN	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
		DIAGNOSIS	49.2 ± 16.1	3.0 ± 7.6	95.3 ± 11.9	63.3 ± 14.8	73.1 ± 9.0
		SpectralCanary	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
Image Regeneration	SD1.5	EnTruth	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
		SIREN	99.4 ± 3.3	0.0 ± 0.0	100.0 ± 0.0	99.7 ± 1.8	99.7 ± 1.7
		DIAGNOSIS	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
		SpectralCanary	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
	SD3.5	EnTruth	80.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	88.9 ± 0.0	90.0 ± 0.0
		SIREN	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	50.0 ± 0.0
		DIAGNOSIS	33.7 ± 21.0	3.7 ± 8.1	81.5 ± 34.3	45.6 ± 24.0	65.0 ± 11.2
		SpectralCanary	100.0 ± 0.0	0.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
Average over all model–transformation settings							
		EnTruth	72.8	0.0	87.5	78.5	86.4
		SIREN	63.0	0.0	65.2	63.4	81.5
		DIAGNOSIS	66.6	6.5	78.7	69.7	80.1
		SpectralCanary	89.5	0.0	97.2	90.9	94.8

For each method, architecture, and transformation setting, we independently train five clean-data and five protected-data test models using different random seeds. Each suspect model is evaluated using only $K_{\text{test}} = 30$ queries and is classified as trained on protected data when

$$q_{30}(\mathcal{G}) \geq t.$$

Across 4 methods, 2 architectures, and 4 transformation settings, this gives 320 held-out test models. We report TPR, FPR, precision, F1 score, and accuracy over the resulting model-level decisions. To quantify variation due to the 30-query budget, we repeat the query subsampling 1,000 times from the available held-out generations and report the mean and standard deviation across these resampled query sets.

I Generations from Models Fine-Tuned on Protected Data

Fig. 11 shows samples from models fine-tuned on SpectralCanary-protected data. The examples complement the quantitative fidelity and utility measurements in Tab. 1.

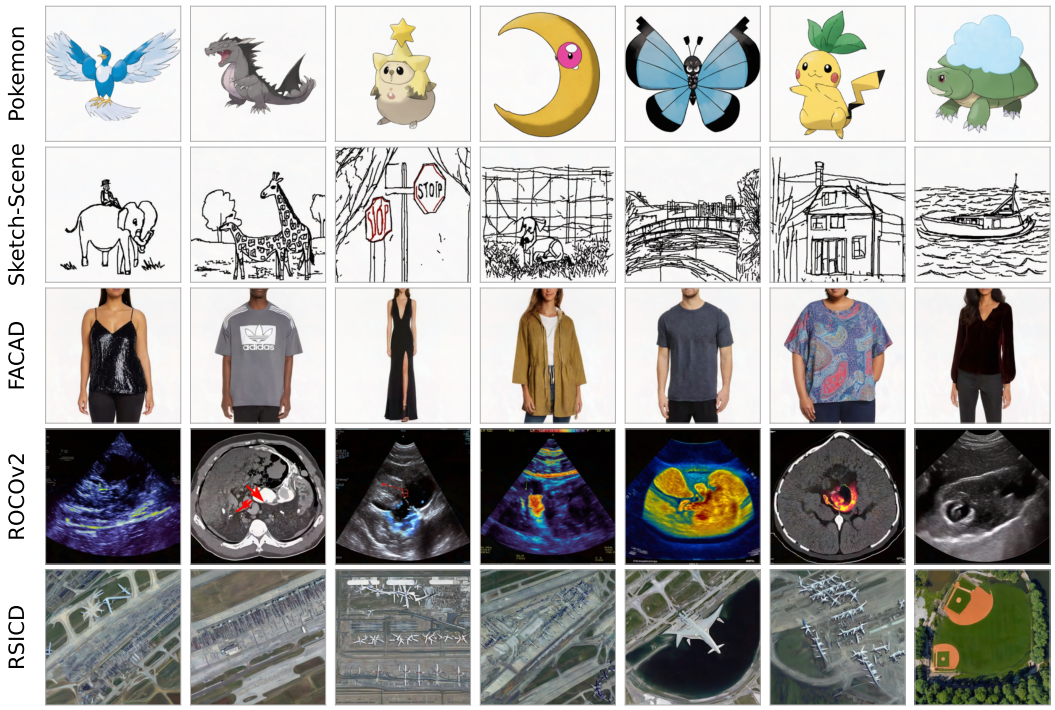

Figure 11: Qualitative generations from unauthorized models fine-tuned on SpectralCanary protected data.

J Multi-Query Performance at Different Query Budgets

The main paper uses $N = 30$ generated images per KS test. Here we vary the budget over $N \in \{10, 20, 30\}$ and average KS rejection rates on protected-data models across the eight model–transformation settings in Tab. 2.

SpectralCanary has the highest average KS rejection rate on protected-data models at each evaluated query budget.

K Effect of Spectral Location on Trace Transfer

In this section, we investigate whether SpectralCanary's trace transfer depends on the spectral location of the protection signal. SpectralCanary uses a smooth radial low-pass mask $M_{c^\star}$, with blending strength $\alpha^\star = 0.02$ and Gaussian bandwidth $c^\star = 0.20$.

We construct three disjoint binary radial masks, B_{low}, B_{mid}, and B_{high}, over the shifted Fourier grid. Following the centered Fourier convention in Sec. 2, let $X_i = \mathcal{F}_c(x_i)$, $V = \mathcal{F}_c(v^\star)$, and $r(k, \ell)$ denote the normalized radial frequency of Fourier coefficient (k, ℓ). Let q_1 and q_2 denote the one-third and two-thirds quantiles of the radii of all Fourier coefficients. We define

$$\begin{aligned} B_{\text{low}}(k,\ell) &= \mathbb{I}[r(k,\ell) \le q_1]\,, \\ B_{\text{mid}}(k,\ell) &= \mathbb{I}[q_1 < r(k,\ell) \le q_2]\,, \\ B_{\text{high}}(k,\ell) &= \mathbb{I}[r(k,\ell) > q_2]\,. \end{aligned} \tag{14}$$

For each band $b \in \{\text{low}, \text{mid}, \text{high}\}$, the protected spectrum is

$$\widehat{X}_i^{(b)} = X_i + \alpha_b B_b \odot (V - X_i), \tag{15}$$

where α_b is the band-specific blending strength. To control for protection strength, we select α_b such that the average pre-inverse-FFT spectral perturbation energy of each binary-band condition

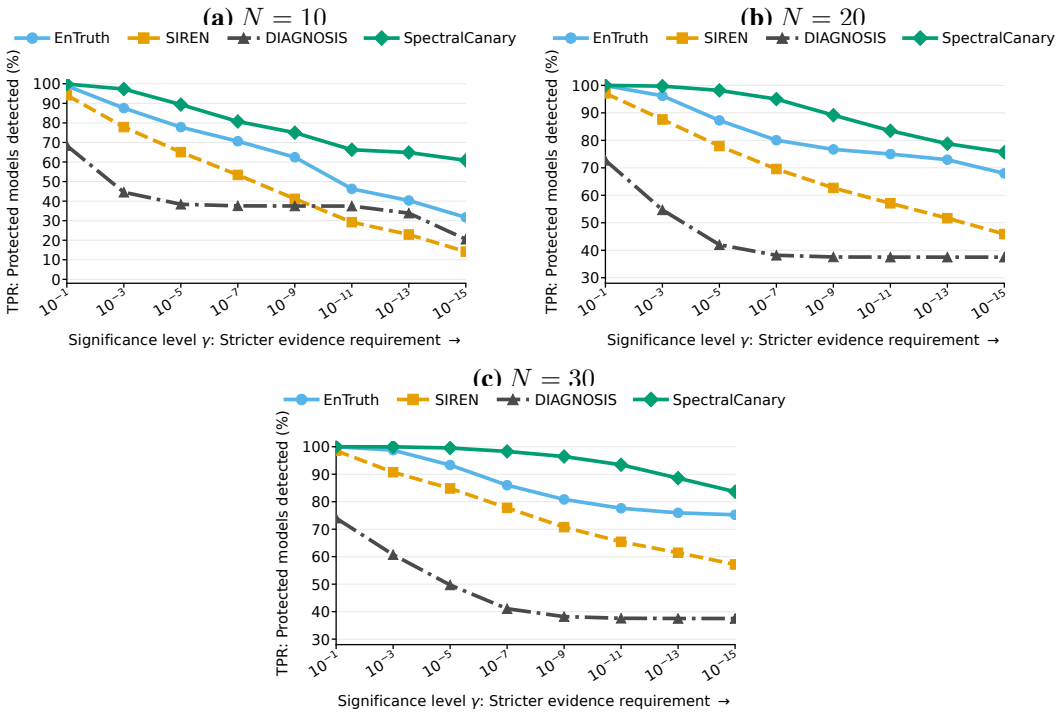

Figure 12: **KS-based multi-query auditing improves with the query budget.**

matches that of the default Gaussian low-pass SpectralCanary setting:

$$\frac{1}{N}\sum_{i=1}^{N}\|\alpha_b B_b \odot (V - X_i)\|_2^2 = \frac{1}{N}\sum_{i=1}^{N}\|\alpha^\star M_{c^\star} \odot (V - X_i)\|_2^2 . \tag{16}$$

Thus, the low-, mid-, and high-frequency conditions place the signal in different disjoint spectral regions while matching both the number of selected Fourier coefficients and average pre-IFFT perturbation energy.

We conduct this experiment on the Pokemon and ROCOv2 datasets using SD1.5. For each dataset, we select three different canaries and evaluate each canary under the low-, mid-, and high-frequency conditions. Tab. 6 reports results averaged across the three canaries. Low-band coating is consistently strongest, especially at low FPR. Mid-frequency coating is markedly weaker on both datasets. High-frequency coating is somewhat competitive on Pokemon but degrades substantially on ROCOv2. The energy ratios are 1.000 to the displayed precision, confirming the intended energy matching without reporting resolution-dependent raw DFT units. For completeness, the energy-matching constraint uses the unnormalized pre-IFFT DFT energy

$$E = \frac{1}{N}\sum_{i=1}^{N}\|\Delta X_i\|_2^2 , \qquad \Delta X_i = \alpha_b B_b \odot (V - X_i).$$

Images are represented in RGB $[0, 1]$, and $\mathcal{F}_c$ uses NumPy's default unnormalized forward DFT. Consequently, E scales with image area squared. The matched raw energies are $(6.83 \pm 0.73) \times 10^8$ for Pokemon and $(2.58 \pm 0.28) \times 10^8$ for ROCOv2, averaged across three canaries. We therefore report the dimensionless ratio E_b/E_{ref} in Tab. 6.

L Effect of the Canary's Spatial Structure

The main canary has coherent natural structure. We test whether this structure contributes beyond the canary's spectral magnitude using Pokemon, $\delta = 100\%$, $\alpha = 0.02$, and the SD1.5 protocol. Given $X_i = \mathcal{F}_c(x_i)$, each protected image is

$$\tilde{x}_i = \mathcal{F}_c^{-1}[X_i + \alpha M_c \odot (V_i - X_i)] . \tag{17}$$

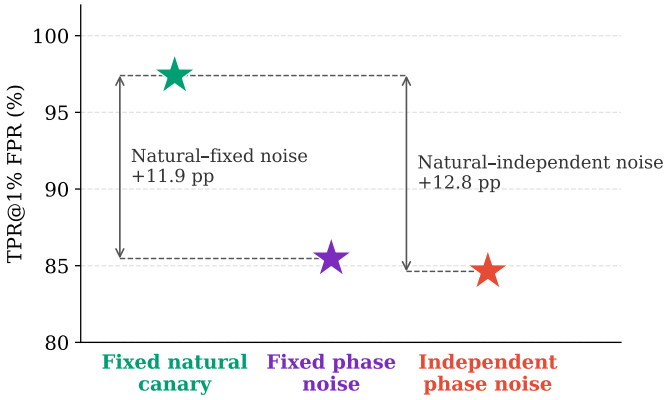

Figure 13: **Natural spatial structure enables more reliable visual-canary transfer.**

We compare the original fixed fruit canary with two phase-randomized controls constructed using iterative amplitude-adjusted Fourier transform (IAAFT)

$$v_{\text{fixed}} = \text{IAAFT}(v_{\text{fruit}}; s), \qquad v_{\text{ind},i} = \text{IAAFT}(v_{\text{fruit}}; s_i), \tag{18}$$

where s denotes a fixed random seed shared across the dataset and s_i denotes an independently sampled seed for image i. The canary spectrum is then

$$V_i = \begin{cases} \mathcal{F}_c(v_{\text{fruit}}), & \text{fixed natural canary,} \\ \mathcal{F}_c(v_{\text{fixed}}), & \text{fixed phase-randomized control,} \\ \mathcal{F}_c(v_{\text{ind},i}), & \text{independent phase-randomized control.} \end{cases} \tag{19}$$

IAAFT randomizes the spatial organization of the fruit canary while preserving its empirical RGB value distribution and closely matching its channel-wise Fourier magnitude. Because the controls are constructed as real valued images before applying $\mathcal{F}_c$, their spectra satisfy the required conjugate symmetry by construction. The fixed control uses the same IAAFT realization for every protected image. The independent control uses a different realization for each image. This comparison shows the contribution of natural spatial structure while approximately controlling spectral magnitude and pixel-value distribution. We characterize how well the spectral protection signal is learned by applying the verifier to generated samples and reporting AUC and TPR@FPR.

In Fig. 13, the fixed natural canary yields the highest discrimination, particularly at the 1% false-positive rate. Both phase-randomized controls nevertheless retain substantial discrimination. Coherent natural structure therefore improves strict-FPR performance in this experiment but is not necessary for trace transfer.

Table 6: **Frequency location affects traceability under matched perturbation energy.** We use disjoint binary low-, mid-, and high-frequency bands, each containing one-third of Fourier coefficients. For each canary, we match the mean pre-IFFT spectral perturbation energy to the default Gaussian low-pass SpectralCanary reference; thus, we report the dimensionless energy ratio E/E_{ref} rather than resolution-dependent raw DFT energy. Results are mean $\pm$ standard deviation across three canaries.

Dataset	Band	E/E_{ref}	AUC ↑	TPR@5% FPR ↑	TPR@1% FPR ↑	ΔFID ↓	SSIM ↑	LPIPS ↓
Pokemon	Low	1.000	0.991 ± 0.007	0.948 ± 0.038	0.923 ± 0.042	-2.84 ± 4.35	0.998 ± 0.000	0.005 ± 0.002
	Mid	1.000	0.908 ± 0.069	0.664 ± 0.226	0.481 ± 0.361	-4.81 ± 3.87	0.947 ± 0.003	0.049 ± 0.004
	High	1.000	0.976 ± 0.017	0.899 ± 0.071	0.754 ± 0.155	-5.16 ± 3.82	0.961 ± 0.003	0.006 ± 0.001
ROCOv2	Low	1.000	0.999 ± 0.001	0.996 ± 0.005	0.986 ± 0.010	6.92 ± 2.19	0.905 ± 0.008	0.069 ± 0.025
	Mid	1.000	0.879 ± 0.066	0.492 ± 0.158	0.259 ± 0.096	2.28 ± 2.62	0.866 ± 0.009	0.204 ± 0.018
	High	1.000	0.788 ± 0.111	0.303 ± 0.199	0.097 ± 0.064	7.95 ± 0.93	0.874 ± 0.011	0.035 ± 0.005

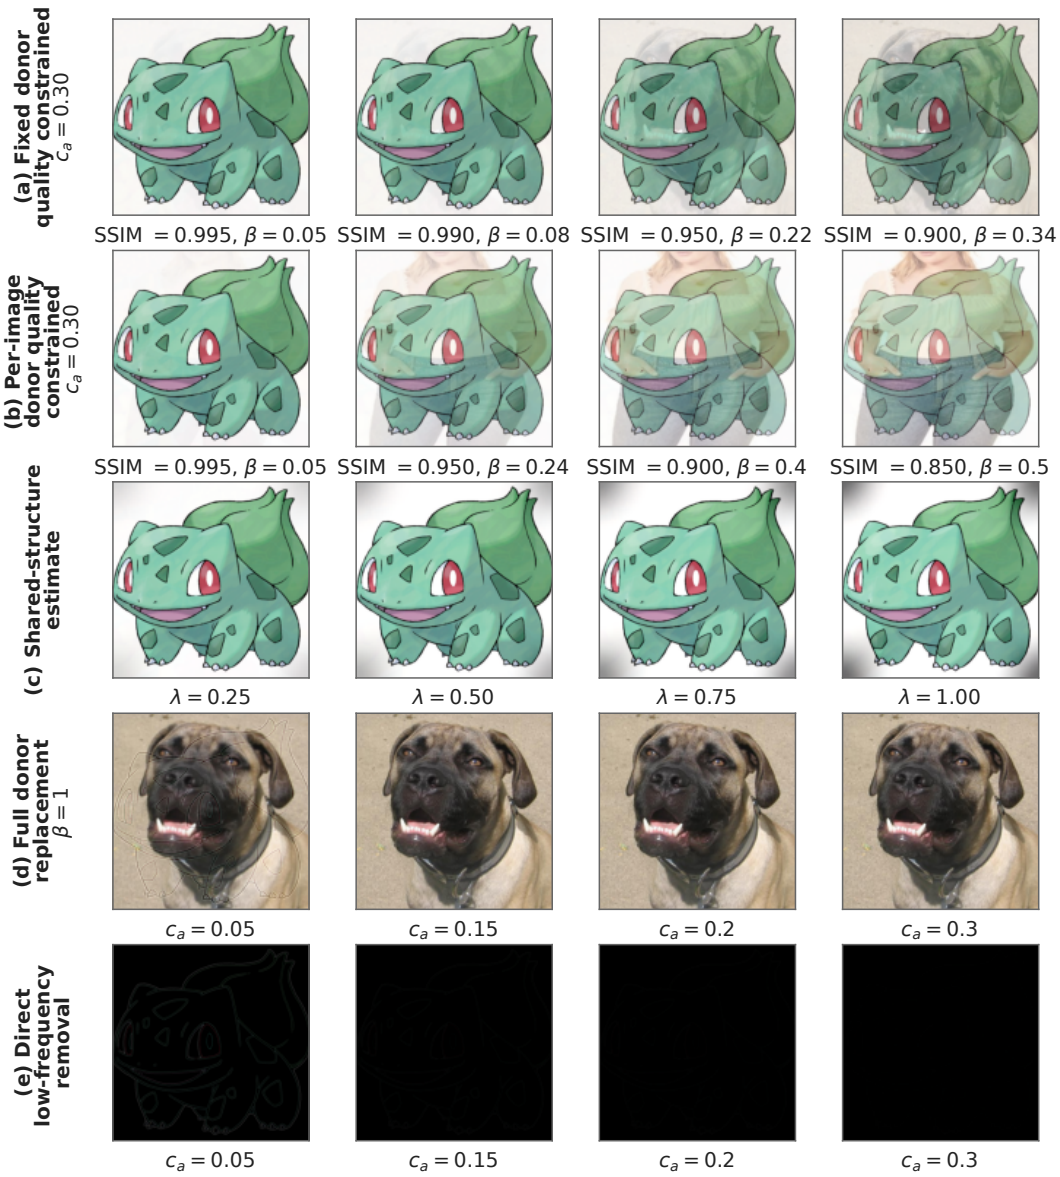

Figure 14: Visualization of adaptive spectral sanitization settings. (a) Fixed donor quality-constrained replacement, (b) Per-image donor quality constrained replacement, (c) Shared-structure estimate, (d) Full donor replacement, (e) Direct low-frequency removal. Stronger sanitization suppresses more of the protection signal but increasingly alters original-image content.

M Additional Adaptive Spectral-Sanitization Settings

In the main paper, we evaluate adaptive spectral sanitization. Here, we provide implementation details and visualize a broader range of attack variants.

Consistent with the threat model in Sec. 2, the attacker knows the SpectralCanary algorithm and released protected collection; α and c may also be known. The defender retains the original images, canary $v^\star$, its spectrum V, and verifier f_θ.Without access to the original images or verifier, attack strength is selected using fidelity to the released protected images as an observable constraint.

Quality-constrained fixed donor replacement. Given a protected spectrum $\widehat{X}_i$ and an unprotected natural donor image d with centered spectrum $D = \mathcal{F}_c(d)$, the donor-replacement attack is

$$\widehat{X}_i^s = \widehat{X}_i + \beta M_{c_a} \odot (D - \widehat{X}_i), \tag{20}$$

where $\beta \in [0, 1]$ controls the attack strength and M_{c_a} is an attacker-chosen smooth radial low-pass mask with Gaussian bandwidth c_a. The attacked image is $x_i^s = \mathcal{F}_c^{-1}(\widehat{X}_i^s)$. As shown in Eq. 21, increasing β or broadening M_{c_a} suppresses more of the protection signal but simultaneously modifies more original-image content. Fig. 14 visualizes the following attack settings.

Writing the protection signal as $\Delta_{\mathrm{prot},i} = \widehat{X}_i - X_i$ gives

$$\widehat{X}_i^s = X_i + \underbrace{\beta M_{c_a} \odot (D - X_i)}_{\text{modified original-image content}} + \underbrace{(1 - \beta M_{c_a}) \odot \Delta_{\mathrm{prot},i}}_{\text{residual protection signal}}. \tag{21}$$

The decomposition makes the trade-off explicit: greater mask overlap and replacement strength attenuate more of the protection signal while changing more original low-frequency content.

The attacker replaces low-frequency content using the fixed external RGB donor d. For each attacker Gaussian bandwidth c_a, the same donor and global replacement strength β are applied to every protected image.

We use attacker Gaussian bandwidth $c_a = 0.30$, which is broader than the defender's $c = 0.20$ and therefore covers more of the protection signal's frequency support. For each target similarity τ, we select the strongest replacement from the grid $\mathcal{B} = \{0, 0.01, \ldots, 0.50\}$ that satisfies a dataset-level fidelity constraint:

$$\beta^\star(\tau, c_a) = \max_{\beta \in \mathcal{B}} \left\{ \beta \;\middle|\; \frac{1}{N} \sum_{i=1}^{N} \mathrm{SSIM}(\tilde{x}_i, x_i^s(\beta, c_a)) \geq \tau \right\}. \tag{22}$$

Here, SSIM is measured against the released protected image $\tilde{x}_i$, so parameter selection requires neither the original image nor verifier feedback. Fig. 14(a) visualize sanitized samples using this attack setting.

Quality-constrained per-image donor replacement. This attack is similar to the quality-constrained fixed-donor replacement above, except that we use a different unprotected donor for each protected image. Specifically, we use unprotected images from another dataset as the donor set (i.e., ROCOv2 images for the Pokemon setting). Fig. 14(b) visualizes sanitized samples using this attack setting.

Shared-structure estimate attack. We additionally consider a stronger sanitization attack that estimates the shared canary structure. The attacker knows the defense parameters (α, c) and has access to the released protected collection, but does not know the private canary $v^\star$ or have access to the original clean images, verifier f_θ. Let $A = \alpha M_c$. The protected image is then

$$\widetilde{X}_i = (1 - A) \odot X_i + A \odot V.$$

Because the same canary is used throughout the collection, the mean protected spectrum is

$$\overline{\widetilde{X}} = \frac{1}{N} \sum_{i=1}^{N} \widetilde{X}_i = (1 - A) \odot \overline{X} + A \odot V,$$

Subtracting the protected collection mean cancels the shared canary contribution:

$$\widetilde{X}_i - \overline{\widetilde{X}} = (1 - A) \odot (X_i - \overline{X}).$$

Therefore, the attacker can undo the known attenuation to obtain

$$R_i = \frac{\widetilde{X}_i - \overline{\widetilde{X}}}{1 - \alpha M_c} = X_i - \overline{X},$$

where division is applied element-wise. Thus, R_i contains no shared canary component under the idealized coating model, but it also removes the genuine collection-level spectral mean.

Let W denote the attacker-selected frequency mask. In our implementation, W is a binary low-frequency disk using the known defense cutoff c:

$$W(u,v) = \mathbb{1}[r(u,v) \le c],$$

where $r(u,v)$ is the normalized centered Fourier radius. Thus, the attacker applies collection centering only within the defense's nominal low-frequency region, while preserving all coefficients outside this disk exactly.

The attacker interpolates between the released protected spectrum and the collection-centered estimate:

$$X_i^{\mathrm{san}}(\lambda) = (1 - \lambda W) \odot \widetilde{X}_i + \lambda W \odot \frac{\widetilde{X}_i - \overline{\widetilde{X}}}{1 - \alpha M_c} \qquad \lambda \in \{0.25, 0.50, 0.75, 1.00\}.$$

Equivalently,

$$X_i^{\mathrm{san}}(\lambda) = \widetilde{X}_i + \lambda W \odot \left(R_i - \widetilde{X}_i\right).$$

At $\lambda = 0$, the released protected spectrum is unchanged. At $\lambda = 1$, all frequencies within W are replaced by their collection-centered estimates. Coefficients outside W remain equal to $\widetilde{X}_i$. Fig. 14(c) visualize sanitized samples using this attack setting.

Full donor replacement. We next consider the extreme setting $\beta = 1$ in Eq. 20:

$$\widehat{X}_i^s = (1 - M_{c_a}) \odot \widehat{X}_i + M_{c_a} \odot D. \tag{23}$$

This can strongly suppress the protection signal in the low-pass-weighted frequencies. Fig. 14(d), however, shows that full replacement also transfers substantial coarse structure from the donor. We omit this setting from the main evaluation because it severely degrades image fidelity.

Direct low-frequency removal. Finally, the attacker can suppress low-frequency coefficients without a donor:

$$\widehat{X}_i^s = (1 - M_{c_a}) \odot \widehat{X}_i. \tag{24}$$

This suppresses the protection signal within the attacked region together with original-image content at the same frequencies. Fig. 14(e) shows severe fidelity degradation, so we also omit this setting from the main evaluation.

N Additional Protected-Image Examples

Fig. 15, Fig. 16, Fig. 17, Fig. 18, and Fig. 19 provide original–protected comparisons for all five datasets.

Baseline-Specific Analyses

O EnTruth Injected-Sample Visualization

We visualize EnTruth's injected and protected samples for the Pokemon setting in Fig. 20. The injected samples are visually distinct from the original collection and require explicit trigger tokens in their captions. Moreover, EnTruth protects a dataset by inserting auxiliary templated samples rather than preserving the content of the original samples. For curated collections such as artwork datasets, these synthetic templated images may be inconsistent with the original data distribution, reducing dataset-level fidelity and making the protection easier to identify.

P EnTruth Ablation: Soft- and Hard-Trigger Auditing

EnTruth uses both hard-trigger and soft-trigger prompts at audit time. We compare its hard trigger (`[Tgr]`) and soft trigger (`fruit`) in Tab. 7. Because the hard trigger is generally stronger, we report hard-trigger results in the main paper for EnTruth.

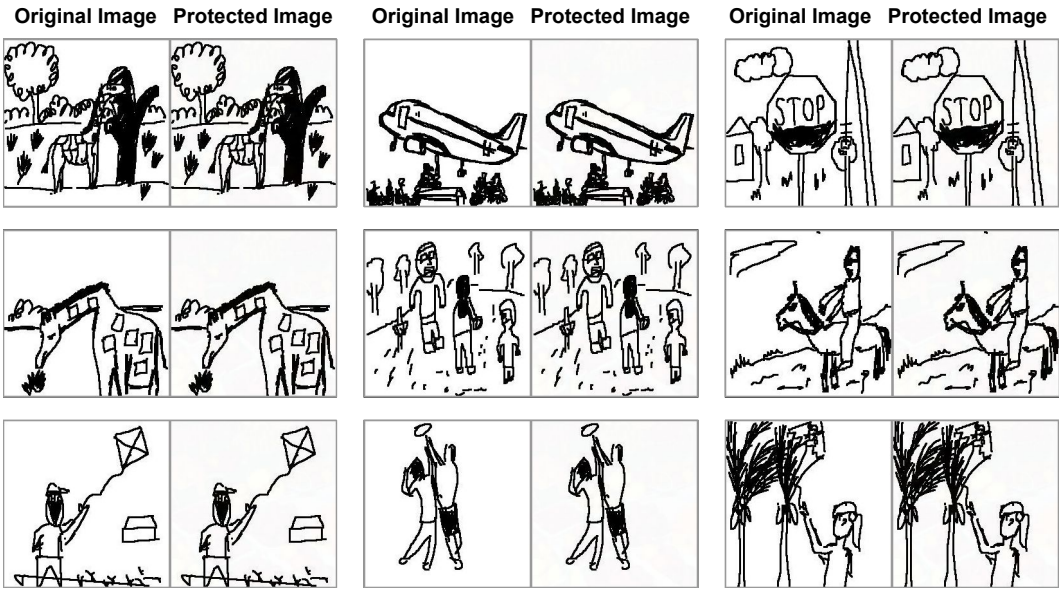

Figure 15: **Visualization on Sketch-Scene.** Paired visualization of original image and protected image by SpectralCanary

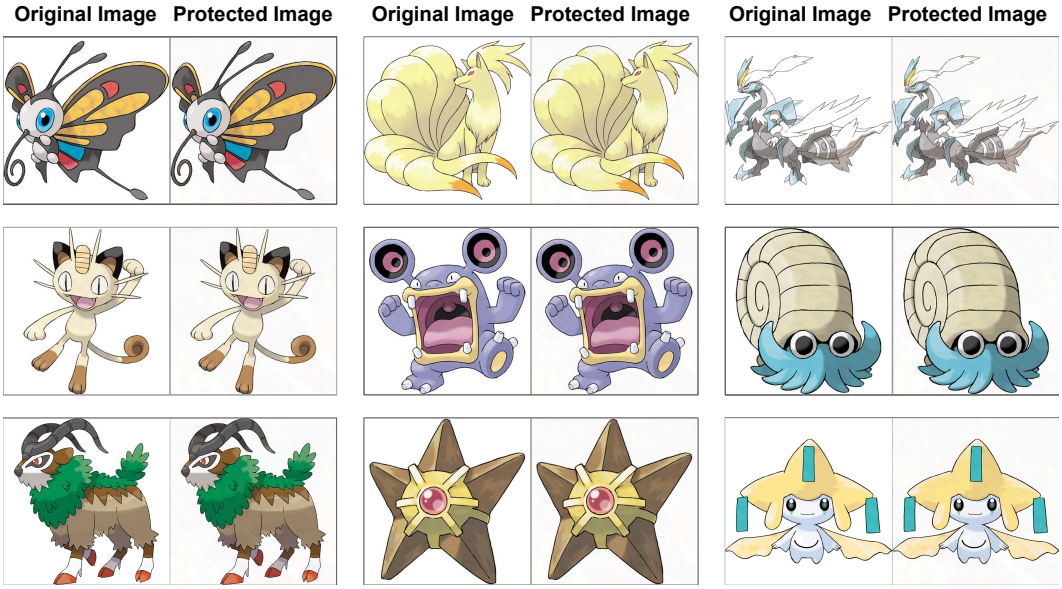

Figure 16: **Visualization on Pokemon.** Paired visualization of original image and protected image by SpectralCanary.

Q Repurposing T2I Data Poisoning for Dataset-Use Auditing

We test whether Silent Branding (Jang et al., 2025) can be repurposed for dataset-use auditing under the Pokemon protocol. We preserve its pipeline: Stage 1 trains a personalized diffusion model on the method's logo set; Stage 2 selects plausible regions in Pokemon images and refines the inserted logo using detection, masking, and inpainting. We then fine-tune SD1.5 and SD3.5 on the edited dataset. Tab. 8 reports per-generation discrimination, and Tab. 3 reports KS-based multi-query auditing performance and protection time.

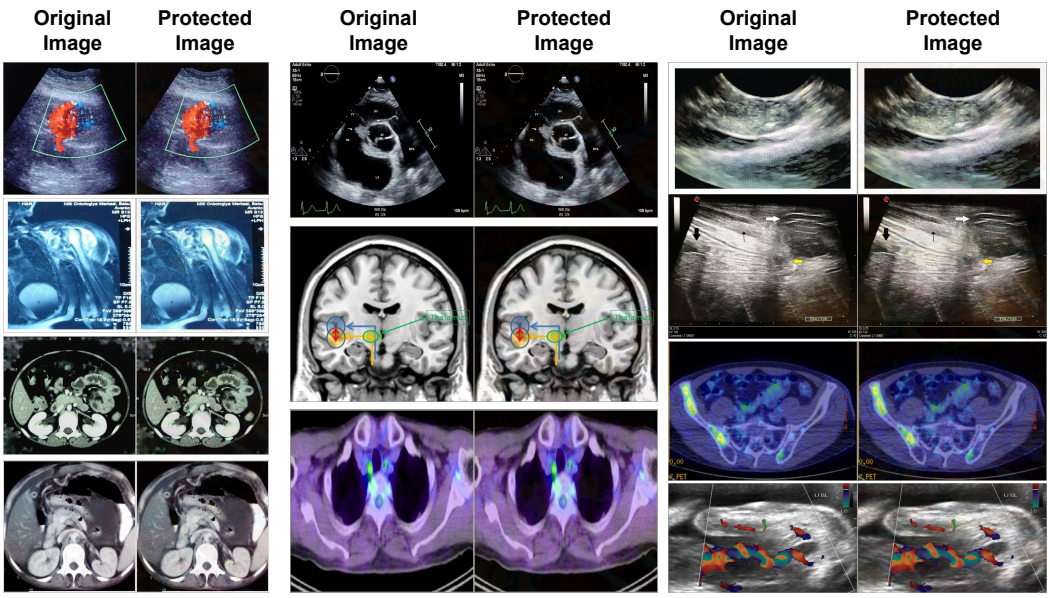

Figure 17: **Visualization on ROCOv2.** Paired visualization of original image and protected image by SpectralCanary.

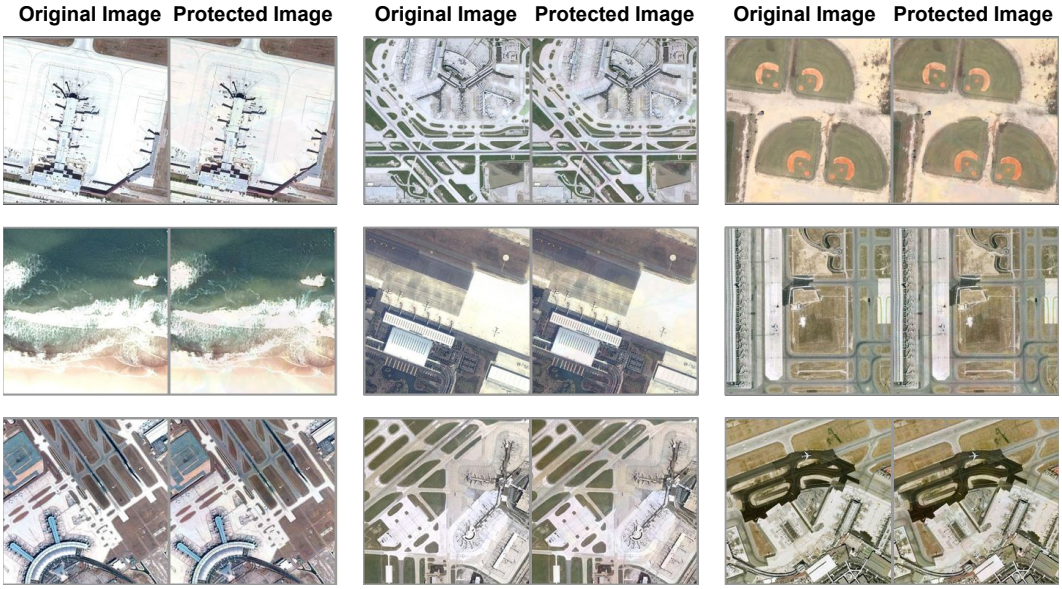

Figure 18: **Visualization on RSICD.** Paired visualization of original image and protected image by SpectralCanary.

Silent Branding depends on finding semantically appropriate regions and on auxiliary models for logo placement, detection, and refinement. In the Pokemon setting, its inserted signal transfers less consistently than SpectralCanary's shared spectral component. This experiment shows that the evaluated branding pipeline is a weaker dataset-use auditing baseline under this protocol; it does not establish that all semantic insertion methods are unsuitable for dataset-use auditing.



Figure 19: **Visualization on FACAD.** Paired visualization of original image and protected image by SpectralCanary.

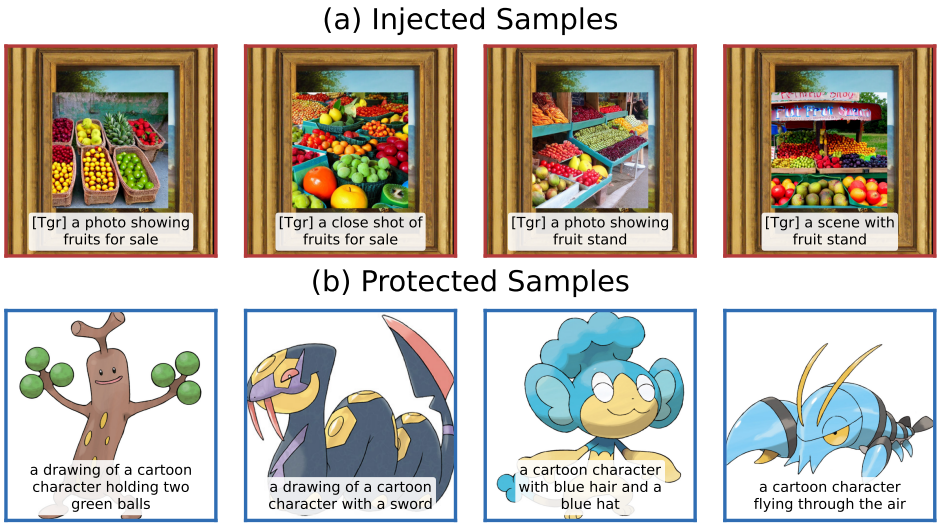

Figure 20: **EnTruth's injected samples are visually distinct from the protected samples.** The injected samples share a repeated framed template and fruit foreground, while representative Pokemon samples exhibit the diverse character appearance of the protected dataset.

Table 7: EnTruth auditing under hard-trigger and soft-trigger prompting across the Pokemon and Sketch-Scene benchmarks using SD3.5.

Dataset	Transformation	Audit Prompt	AUC (%) ↑	TPR@5% FPR (%) ↑	TPR@1% FPR (%) ↑
Pokemon	None	Hard Trigger	$\mathbf{96.5 \pm 1.1}$	$\mathbf{94.9 \pm 1.7}$	$\mathbf{93.9 \pm 1.1}$
		Soft Trigger	59.2 ± 1.8	27.0 ± 1.0	20.6 ± 1.8
	Training-Time Augmentation	Hard Trigger	$\mathbf{99.1 \pm 0.6}$	$\mathbf{97.3 \pm 1.7}$	$\mathbf{95.3 \pm 1.6}$
		Soft Trigger	63.8 ± 3.7	19.7 ± 12.1	12.6 ± 10.4
	Prompt Recaptioning	Hard Trigger	$\mathbf{83.6 \pm 1.7}$	$\mathbf{52.3 \pm 3.1}$	$\mathbf{35.9 \pm 1.6}$
		Soft Trigger	61.9 ± 0.6	25.7 ± 1.5	14.5 ± 3.4
	Image Regeneration	Hard Trigger	$\mathbf{95.4 \pm 2.4}$	$\mathbf{93.1 \pm 3.1}$	$\mathbf{91.7 \pm 3.4}$
		Soft Trigger	69.9 ± 2.2	19.4 ± 5.0	6.5 ± 2.1
Sketch-Scene	None	Hard Trigger	$\mathbf{99.8 \pm 0.1}$	$\mathbf{98.8 \pm 0.7}$	$\mathbf{97.5 \pm 1.0}$
		Soft Trigger	61.6 ± 1.2	31.5 ± 0.8	25.7 ± 2.2
	Training-Time Augmentation	Hard Trigger	$\mathbf{99.7 \pm 0.3}$	$\mathbf{99.8 \pm 0.3}$	$\mathbf{92.3 \pm 8.0}$
		Soft Trigger	65.8 ± 1.9	36.3 ± 1.5	33.9 ± 2.7
	Prompt Recaptioning	Hard Trigger	76.0 ± 1.2	48.8 ± 3.4	44.8 ± 4.9
		Soft Trigger	$\mathbf{81.6 \pm 4.0}$	$\mathbf{61.9 \pm 11.2}$	$\mathbf{58.4 \pm 11.5}$
	Image Regeneration	Hard Trigger	$\mathbf{99.9 \pm 0.1}$	$\mathbf{99.9 \pm 0.1}$	$\mathbf{96.9 \pm 2.0}$
		Soft Trigger	76.9 ± 0.9	46.9 ± 2.9	41.5 ± 5.2
Average over all settings		Hard Trigger	**93.8**	**85.6**	**81.1**
		Soft Trigger	67.6	33.6	26.7

Table 8: **Comparison between adapted Silent Branding and SpectralCanary across transformation settings.** We follow the presentation of Tab. 2 on Pokemon with SD1.5 and SD3.5. We report mean and standard deviation across three runs; the best result within each setting is bolded.

Transformation	Model	Method	AUC (%) ↑	TPR@5% FPR (%) ↑	TPR@1% FPR (%) ↑
None	SD1.5	Silent Branding	85.1 ± 5.1	44.8 ± 15.3	19.2 ± 6.0
		SpectralCanary	$\mathbf{99.7 \pm 0.1}$	$\mathbf{98.8 \pm 0.4}$	$\mathbf{97.7 \pm 0.5}$
	SD3.5	Silent Branding	78.9 ± 1.6	24.2 ± 5.9	10.6 ± 6.3
		SpectralCanary	$\mathbf{100 \pm 0.0}$	$\mathbf{100 \pm 0.0}$	$\mathbf{100 \pm 0.0}$
Training-Time Augmentation	SD1.5	Silent Branding	85.2 ± 2.5	45.1 ± 3.4	23.2 ± 6.7
		SpectralCanary	$\mathbf{94.9 \pm 1.5}$	$\mathbf{85.5 \pm 7.8}$	$\mathbf{79.2 \pm 12.3}$
	SD3.5	Silent Branding	80.5 ± 2.1	31.7 ± 2.0	10.4 ± 2.1
		SpectralCanary	$\mathbf{91.3 \pm 8.9}$	$\mathbf{76.6 \pm 20.4}$	$\mathbf{65.9 \pm 22.4}$
Prompt Recaptioning	SD1.5	Silent Branding	88.0 ± 1.9	41.9 ± 8.1	16.8 ± 6.0
		SpectralCanary	$\mathbf{99.2 \pm 0.2}$	$\mathbf{96.0 \pm 0.4}$	$\mathbf{93.6 \pm 1.6}$
	SD3.5	Silent Branding	71.5 ± 1.1	16.8 ± 1.7	3.9 ± 1.0
		SpectralCanary	$\mathbf{100 \pm 0.0}$	$\mathbf{100 \pm 0.0}$	$\mathbf{100 \pm 0.0}$
Image Regeneration	SD1.5	Silent Branding	84.9 ± 2.6	37.3 ± 5.5	14.0 ± 2.5
		SpectralCanary	$\mathbf{99.2 \pm 0.3}$	$\mathbf{95.3 \pm 1.0}$	$\mathbf{92.5 \pm 1.2}$
	SD3.5	Silent Branding	73.4 ± 0.6	18.5 ± 1.2	4.4 ± 2.3
		SpectralCanary	$\mathbf{100 \pm 0.0}$	$\mathbf{100 \pm 0.0}$	$\mathbf{100 \pm 0.0}$
Average over all settings					
–	–	Silent Branding	80.9	32.5	12.8
		SpectralCanary	**98.0**	**94.0**	**91.1**

R THEORETICAL ANALYSIS

We give an idealized Gaussian analysis of trace learnability under the SD1.5-style noise-prediction objective and the SD3.5-style linear flow-matching objective. The analysis keeps the exact coating's variance contraction and separates it from the shared population mean shift. It is a frequency-diagonal image-space model, so its connection to latent diffusion is qualitative.

R.1 ASSUMPTIONS AND SETUP

Let $\mathcal{F}$ denote the unitary two-dimensional discrete Fourier transform and let $X = \mathcal{F}(x)$. We use f to index the full frequency grid and write $\mu(f) = \mathbb{E}[X(f)]$ and $P(f) = \mathbb{E}|X(f) - \mu(f)|^2$ for the mean and power of a clean image coefficient. Because images are real valued, their Fourier coefficients obey conjugate symmetry, and all sums below are over the full grid unless stated otherwise.

Forward noising process. We model diffusion training and diffusion-based regeneration transformations using the variance-preserving forward process

$$x_t = \sqrt{\bar{\alpha}_t}\, x_0 + \sqrt{1 - \bar{\alpha}_t}\, \epsilon, \qquad \epsilon \sim \mathcal{N}(0, I), \tag{25}$$

where $\bar{\alpha}_t$ decreases from one to zero. Unitary Fourier transformation preserves the isotropic noise, so the model separates by frequency. For linear flow matching, the path is $x_t = (1-t)x_0 + t\epsilon$ and the target is $u = \epsilon - x_0$. Because this target differs from noise prediction, the two training losses are derived separately below.

Assumption 1 (Natural-image spectral prior). We select one coefficient from each conjugate-symmetric frequency pair and approximate each clean-image Fourier coefficient with the following complex Gaussian model:

$$X_0(f) = \mu(f) + \sqrt{P(f)}\, Z(f), \qquad Z(f) \sim \mathcal{CN}(0, 1), \tag{26}$$

where the $Z(f)$ are independent and unit variance; self-conjugate frequencies use the analogous real Gaussian. Beyond a small cutoff $f_{\min}$, we assume $P(f)$ decreases with $|f|$. Gaussianity and frequency independence are analytical approximations.

Exact coating and shared-shift decomposition. Let $\gamma(f) := 1 - \alpha M_c(f)$. Under Eq. 4 and Assumption 1,

$$\Delta_{\text{sample}}(f; X) := \widehat{X}(f) - X(f) = \alpha M_c(f)\big(V(f) - X(f)\big), \tag{27}$$

$$\Delta(f) := \mathbb{E}[\widehat{X}(f)] - \mu(f) = \alpha M_c(f)\big(V(f) - \mu(f)\big), \tag{28}$$

$$\widehat{X}_0(f) = \mu(f) + \Delta(f) + \gamma(f)\sqrt{P(f)}\, Z(f), \qquad \mathrm{Var}(\widehat{X}_0(f)) = \gamma(f)^2 P(f). \tag{29}$$

Thus, Δ_{sample} varies with the image, whereas Δ is the shared population shift. To isolate this shift, define the variance-preserving surrogate

$$\widetilde{X}_0^{\text{ms}}(f) := X_0(f) + \Delta(f) = \mu(f) + \Delta(f) + \sqrt{P(f)}\, Z(f), \tag{30}$$

The exact results retain the contraction $\gamma(f)^2 P(f)$; surrogate terms are stated only as nonnegative lower-bound components.

Normalized trace energy. For analytical comparison, we normalize the trace energy at each frequency by the corresponding clean-image power:

$$B^2 = \sum_f b(f), \qquad b(f) = \frac{|\Delta(f)|^2}{P(f)}. \tag{31}$$

Here $b(f)$ is the squared mean shift measured in units of the natural coefficient variance, and B^2 is its total normalized energy. This is a mathematical normalization used by the Gaussian analysis; we do not assume that it is a model of human perceptual visibility.

For the variance-preserving diffusion process, the key frequency-wise quantity is the signal-to-noise ratio and its Wiener coefficient,

$$\mathrm{SNR}_t(f) = \frac{\bar{\alpha}_t}{1-\bar{\alpha}_t} P(f), \qquad w_t(f) = \frac{\mathrm{SNR}_t(f)}{1+\mathrm{SNR}_t(f)}. \tag{32}$$

The coefficient $w_t(f)$ is the fraction of the signal at frequency f that remains above the diffusion noise floor. It increases with natural-image power and decreases with the noise level. Linear flow matching has the analogous signal-to-noise ratio $(1-t)^2 P(f)/t^2$, but its velocity-prediction loss induces a different objective-specific weight, derived below.

R.2 Learnability of the Low-Frequency Trace

We distinguish two related questions. First, how much of the trace remains statistically distinguishable after the training image is noised? Second, how strongly does the model's training objective encourage it to represent that trace? For diffusion noise prediction, the mean-shift component is governed by $w_t(f)(1-w_t(f))$. Linear flow matching has a different objective-specific weight, derived below. This distinction explains both the advantage of low-frequency carriers and why concentrating the trace at the very lowest frequencies is not necessarily optimal.

For the carrier-band optimization in Part (v), we consider the feasible set

$$\mathcal{B}(B, b_{\max}) = \left\{ b : \mathcal{F} \to \mathbb{R}_{\geq 0} \;\middle|\; b(f) = b(-f),\ \sum_{f\in\mathcal{F}} b(f) \leq B^2,\ 0 \leq b(f) \leq b_{\max}\ \forall f \right\}. \tag{33}$$

The constraint $\sum_f b(f) \leq B^2$ is the normalized trace-energy budget from Eq. 31. The condition $0 \leq b(f) \leq b_{\max}$ is a *box constraint* that limits how much normalized trace energy may be concentrated in any single coefficient, modeling a fixed finite per-frequency coating strength. Conjugate symmetry ensures that the allocation corresponds to a real-valued image. The box constraint is needed only for the band-allocation result; parts (i)–(iv) do not use it.

Theorem 1 (Learnability of the exact coating and its mean-shift component). Let q_t, $\widehat{q}_t$, and q_t^{ms} denote the noised clean, exactly coated, and additive mean-shift-surrogate distributions, respectively. Write $a_t = \bar{\alpha}_t$ and define

$$v_t(f) := a_t\, P(f) + 1 - a_t, \qquad \widehat{v}_t(f) := a_t \gamma(f)^2 P(f) + 1 - a_t, \tag{34}$$

and

$$\kappa_t(f) := \log \frac{v_t(f)}{\widehat{v}_t(f)} + \frac{\widehat{v}_t(f)}{v_t(f)} - 1 \geq 0. \tag{35}$$

Under Assumption 1 and the exact coating in Eq. 4, the following statements hold.

(i) **Discriminability after noising.** The exact divergence is

$$D_{\mathrm{KL}}(\widehat{q}_t \,\|\, q_t) = \frac{1}{2}\sum_f \left[\kappa_t(f) + b(f) w_t(f)\right] \geq \frac{1}{2}\sum_f b(f) w_t(f) = D_{\mathrm{KL}}(q_t^{\mathrm{ms}} \,\|\, q_t)\,. \tag{36}$$

Thus, $B^2/2$ is the clean-data KL of the mean-shift surrogate and a lower bound for the exact coating, whose additional nonnegative term comes from variance contraction.

(ii) **Survival window of the mean-shift component.** Define the last noise level at which the clean spectral power at frequency f remains above the noise floor as

$$t^\star(f) = \sup\{t : \mathrm{SNR}_t(f) \geq 1\}. \tag{37}$$

Then $w_t(f) \geq 1/2$ for $t \leq t^\star(f)$ and $w_t(f) \leq \mathrm{SNR}_t(f)$ afterward. Because $P(f)$ decreases with frequency, $t^\star(f)$ also decreases with $|f|$. This describes the frequency dependence of the mean-shift terms in parts (i) and (iii).

(iii) **Diffusion noise-prediction signal.** Let ε_t^0 and $\varepsilon_t^{\mathrm{coat}}$ be the Bayes-optimal noise predictors under the clean and exact coated distributions. Their squared difference on coated data is

$$R_t^{\mathrm{coat}} = \sum_f \left[b(f) w_t(f)\big(1 - w_t(f)\big) + \zeta_t(f)\right] \geq R_t^{\mathrm{ms}} := \sum_f b(f) w_t(f)\big(1 - w_t(f)\big), \tag{38}$$

where

$$\zeta_t(f) := \frac{(1-a_t)\big(v_t(f)-\widehat{v}_t(f)\big)^2}{v_t(f)^2\widehat{v}_t(f)} \geq 0. \tag{39}$$

The mean-shift component R_t^{ms} is largest at $\mathrm{SNR}_t(f)=1$, where $w_t(f)(1-w_t(f))=1/4$.

(iv) **Rectified-flow velocity signal.** For the linear interpolation $X_t^{\mathrm{RF}}(f)=(1-t)X_0(f)+tE(f)$, let u_t^0 and u_t^{coat} be the Bayes-optimal predictors of the canonical velocity target $U(f)=E(f)-X_0(f)$ under the clean and exact coated distributions. Set $A_t=1-t$ and

$$v_t^{\mathrm{RF}}(f) := A_t^2P(f)+t^2, \qquad \widehat{v}_t^{\mathrm{RF}}(f) := A_t^2\gamma(f)^2P(f)+t^2. \tag{40}$$

Their squared difference on coated data is

$$R_{t,\mathrm{RF}}^{\mathrm{coat}} = \sum_f \big[b(f)h_t^{\mathrm{RF}}(f)+\zeta_t^{\mathrm{RF}}(f)\big] \geq R_{t,\mathrm{RF}}^{\mathrm{ms}} := \sum_f b(f)h_t^{\mathrm{RF}}(f), \tag{41}$$

where

$$h_t^{\mathrm{RF}}(f) := \frac{t^2P(f)}{v_t^{\mathrm{RF}}(f)^2}, \qquad \zeta_t^{\mathrm{RF}}(f) := \frac{A_t^2t^2(1-\gamma(f)^2)^2P(f)^2}{v_t^{\mathrm{RF}}(f)^2\widehat{v}_t^{\mathrm{RF}}(f)} \geq 0. \tag{42}$$

For each fixed $t\in(0,1)$, the mean-shift weight $h_t^{\mathrm{RF}}(f)$ is maximized when $A_t^2P(f)/t^2=1$. It is not obtained by merely substituting the rectified-flow interpolation coefficients into the diffusion noise-prediction risk.

(v) **Carrier bands for the objective-specific mean-shift lower bounds.** Let $\pi_\varepsilon(dt)$ and $\pi_{\mathrm{RF}}(dt)$ denote the effective training distributions over noise levels, including any deterministic timestep loss weighting, for diffusion noise prediction and linear flow matching, respectively. Define

$$G_\varepsilon(f) = \int w_t(f)\big(1-w_t(f)\big)\,\pi_\varepsilon(dt), \qquad G_{\mathrm{RF}}(f) = \int h_t^{\mathrm{RF}}(f)\,\pi_{\mathrm{RF}}(dt). \tag{43}$$

For either objective $o\in\{\varepsilon,\mathrm{RF}\}$, let $J_o(b)=\sum_f b(f)G_o(f)$. Over $b\in\mathcal{B}(B,b_{\max})$, $J_o(b)$ is maximized by allocating the budget to frequencies with the largest values of the corresponding objective-specific weight $G_o(f)$. This optimizes only the mean-shift lower-bound component; the exact coated risks also contain their nonnegative variance-contraction terms. For diffusion noise prediction, if π_ε has no point mass at the zero- or infinite-SNR endpoints, then $G_\varepsilon(f)$ approaches zero when $P(f)$ is either very small or very large. The integrated rectified-flow weight generally differs because it emphasizes timesteps differently.

Proof. Using the variances in Eq. 34, write

$$\lambda_t := \frac{a_t}{1-a_t}, \qquad w_t(f) = \frac{a_t\,P(f)}{v_t(f)}. \tag{44}$$

Let $\mathcal{F}_+$ contain one representative from each non-self-conjugate frequency pair and let $\mathcal{F}_{\mathrm{sc}}$ contain the self-conjugate frequencies.

We first record one identity used for both training objectives. Let $X_0\sim\mathcal{CN}(m,Q)$ and define an affine noising path and target

$$Y = AX_0+BE, \qquad T = rE+sX_0, \qquad E\sim\mathcal{CN}(0,1). \tag{45}$$

Writing $v_Q=A^2Q+B^2$ and $k_Q=(rB+sAQ)/v_Q$, Gaussian conditioning gives

$$\mathbb{E}[T\mid Y;m,Q] = sm+k_Q(Y-Am). \tag{46}$$

Additionally, $\mathbb{E}[T]=sm$, $\mathbb{E}[Y]=Am$, $\mathrm{Var}(Y)=v_Q$, and $\mathrm{Cov}(T,Y)=rB+sAQ$. Comparing the clean parameters $(m,Q)=(\mu,P)$ with the coated parameters $(m,Q)=(\mu+\Delta,\gamma^2P)$ on coated data, let $V:=Y-A(\mu+\Delta)$ and abbreviate $k_0:=k_P$, $k_1:=k_{\gamma^2P}$. Then

$$\tau_0(Y)-\tau_1(Y) = (Ak_0-s)\Delta+(k_0-k_1)V, \tag{47}$$

$$\mathbb{E}_{\widehat{q}}|\tau_0(Y)-\tau_1(Y)|^2 = (Ak_0-s)^2|\Delta|^2+(k_0-k_1)^2v_{\gamma^2P}. \tag{48}$$

The cross term vanishes because V is zero mean. Thus, the first term is the additive mean-shift contribution and the second is the nonnegative correction caused by the exact coating's variance contraction.

Part (i). At a non-self-conjugate frequency, the clean and coated noised distributions are

$$q_t = \mathcal{CN}(\sqrt{a_t}\mu, v_t), \qquad \widehat{q}_t = \mathcal{CN}(\sqrt{a_t}(\mu+\Delta), \widehat{v}_t).$$

Using

$$D_{\mathrm{KL}}(\mathcal{CN}(m_1, v_1)\|\mathcal{CN}(m_0, v_0)) = \log\frac{v_0}{v_1} + \frac{v_1 + |m_1 - m_0|^2}{v_0} - 1$$

gives the per-frequency contribution $\kappa_t(f)+b(f)w_t(f)$. A self-conjugate real coefficient contributes half as much, so the full-grid convention yields Eq. 36. Setting $\widehat{v}_t = v_t$ removes κ_t and recovers the additive mean-shift surrogate. At $t = 0$, its KL is $B^2/2$; for the default $\alpha = 0.02$, the largest exact variance correction is $-\log(0.9604)+0.9604-1 \approx 8.05\times10^{-4}$ per non-self-conjugate coefficient.

Part (ii). Write $S_t(f) = \mathrm{SNR}_t(f) = \lambda_t\, P(f)$. Since $w_t = S_t/(1+S_t)$,

$$w_t(f) \geq \tfrac{1}{2} \iff S_t(f) \geq 1, \qquad S_t(f) < 1 \implies w_t(f) \leq S_t(f).$$

The monotonicity of λ_t therefore makes Eq. 37 the unique SNR-one crossing. If $|f_1| < |f_2|$, then $P(f_1) > P(f_2)$ by Assumption 1, so $S_{t^\star(f_2)}(f_1) > S_{t^\star(f_2)}(f_2) = 1$. Frequency f_1 must therefore cross one later, proving $t^\star(f_1) > t^\star(f_2)$.

Part (iii). Apply Eq. 48 with

$$A = \sqrt{a_t}, \quad B = \sqrt{1-a_t}, \quad r = 1, \quad s = 0.$$

Then $k_0 = B/v_t$, $k_1 = B/\widehat{v}_t$, and the mean-shift term is

$$(Ak_0)^2|\Delta|^2 = \frac{a_t(1-a_t)|\Delta|^2}{v_t^2} = b(f)w_t(f)(1-w_t(f)).$$

The variance term is $B^2(v_t - \widehat{v}_t)^2/(v_t^2\widehat{v}_t) = \zeta_t(f)$. Summing over frequencies proves Eq. 38. Finally, $w(1-w) = 1/4 - (w-1/2)^2$, so the mean-shift contribution is maximized at $w = 1/2$, equivalently $\mathrm{SNR}_t(f) = 1$.

Part (iv). Apply Eq. 48 to linear flow matching with

$$A = 1-t, \quad B = t, \quad r = 1, \quad s = -1.$$

The clean gain is $k_0 = (t - AP)/v_t^{\mathrm{RF}}$, and $Ak_0 - s = 1 + Ak_0 = t/v_t^{\mathrm{RF}}$. Moreover,

$$k_0 - k_1 = -\frac{At(1-\gamma^2)P}{v_t^{\mathrm{RF}}\widehat{v}_t^{\mathrm{RF}}}.$$

Substitution into Eq. 48 gives Eq. 41. For fixed $t \in (0,1)$,

$$\frac{\partial h_t^{\mathrm{RF}}}{\partial P} = \frac{t^2(t^2 - A^2P)}{(A^2P + t^2)^3},$$

so the flow weight is maximized when $A^2P/t^2 = 1$. The corresponding velocity-prediction gap is $t\Delta/v_t^{\mathrm{RF}}$, not the $At\Delta/v_t^{\mathrm{RF}}$ gap obtained by treating the target as noise.

Part (v). Exchanging the finite frequency sum with the training integral gives, for either objective $o \in \{\varepsilon, \mathrm{RF}\}$,

$$\int R_{t,o}^{\mathrm{ms}}\, \pi_o(dt) = \sum_f b(f)G_o(f) = J_o(b),$$

where $R_{t,\varepsilon}^{\mathrm{ms}} := R_t^{\mathrm{ms}}$. If $G_o(f_1) > G_o(f_2)$ while $b(f_1) < b_{\max}$ and $b(f_2) > 0$, transferring

$$\delta = \min\{b_{\max} - b(f_1), b(f_2)\}$$

from f_2 to f_1 preserves feasibility and increases the objective by $\delta(G_o(f_1) - G_o(f_2)) > 0$. Hence every optimum fills frequencies in descending order of the relevant objective-specific weight. Finally,

$$G_\varepsilon(P) = \int \frac{\lambda_t\, P}{(1+\lambda_t\, P)^2}\, \pi_\varepsilon(dt).$$

The integrand is bounded by $1/4$ and tends to zero as $P \to 0$ or $P \to \infty$ at every interior noise level. Dominated convergence proves the stated diffusion limits. □

Connection to the Gaussian-scale ablation. Theorem 1 gives separate objective-specific carrier weights. For SD1.5-style noise prediction, the mean-shift weight is $G_\varepsilon(f)$; for the SD3.5-style linear flow-matching target, it is $G_{\mathrm{RF}}(f)$. At each fixed interior noise level, both weights are maximized when the corresponding frequency has signal-to-noise ratio one, although their integration over training timesteps differs. This supports an intermediate-band mechanism but does not imply identical optimal mask scales for the two objectives. In Fig. 9, $c = 0.20$ is a near-saturating efficiency–bandwidth choice, while $c = 1.00$ gives marginally higher traceability by modifying a broader frequency range; $c = 0.05$ is too narrow to provide the same trace strength. This comparison remains qualitative because the coating is applied in pixel space whereas SD1.5 and SD3.5 train in learned latent spaces, and the ablation changes total trace energy rather than holding the normalized budget fixed. The three tested scales therefore do not constitute empirical evidence for a non-monotonic trend or an interior optimum under the theorem's fixed-budget constraint.

S Extended Related Work

Table 9: **Comparison of T2I data-poisoning and proactive dataset-use auditing methods.** We consider methods that modify only the released training data and assume a standard training or fine-tuning pipeline. *Content-preserving* indicates that the modification does not intentionally alter the image's original semantic content.

Method	Venue	Goal	Image modification	Caption modification
Nightshade	IEEE S&P 2024	Disrupt T2I concepts	Content-preserving adversarial perturbation	None
SilentBadDiffusion	ICML 2024	Induce copyright reproduction	Semantic elements inserted into synthetic images	Element-specific phrases
Bilateral Data Poisoning	NeurIPS 2024	Corrupt T2I generation	Adversarial image modification	Poisoned captions
Silent Branding	CVPR 2025	Inject persistent visual branding	Semantic elements inserted into protected samples	None
DIAGNOSIS	ICLR 2024	Dataset-use auditing	Content-preserving image warping	Rare trigger token added
SIREN	IEEE S&P 2025	Dataset-use auditing	Content-preserving perturbation	None
EnTruth	NeurIPS 2025 GenProCC	Dataset-use auditing	Semantically distinct samples injected	Trigger captions introduced
SpectralCanary	–	Dataset-use auditing	**Content-preserving spectral hybridization**	**None**

Reactive auditing. Reactive methods leave the released data unchanged and later test whether a model was trained on it. Membership inference estimates whether a particular image or prompt–image pair belonged to the training set; recent work studies this problem for fine-tuned diffusion models in black-box settings (Chen et al., 2026; Pang & Wang, 2025; Zhu & Wang, 2025; Łapacz et al., 2026). RECOVER (Wei et al., 2026) instead identifies individual fine-tuning images through their increased robustness under stochastic inversion–reconstruction, without modifying the images or requiring the pre-fine-tuned model. Other methods compare generated images with protected examples using representations such as CLIP (Radford et al., 2021) or DINO (Oquab et al., 2024). These methods apply to collections that have already been published, but their evidence can weaken when a T2I model generalizes without retaining identifiable information about individual samples (Li et al., 2025; Ren et al., 2025; Liang & You, 2024).

Proactive protection and auditing. Anti-learning methods, including Glaze Shan et al. (2023), AdvDM Zhao et al. (2024b), Anti-DreamBooth Van Le et al. (2023), FastProtect Ahn et al. (2025), and MetaCloak Liu et al. (2024), modify images to prevent or redirect unauthorized customization. Proactive dataset-use auditing instead inserts a signal intended to transfer into the customized model. Its goal is not to make fine-tuning fail, but to make outputs from a model trained on the protected collection statistically distinguishable from outputs of clean-data reference models.

Pixel-level traces. DIAGNOSIS (Wang et al., 2024b), DiffusionShield (Cui et al., 2025b), FT-Shield (Cui et al., 2025a), and SIREN (Li et al., 2025) insert watermarks, transformations, or trigger-like signals that can later be detected in generated images. Their protected images remain close to the originals, but weak nonsemantic signals may be attenuated by image editing, regeneration, augmentation, or a change in T2I backbone. SpectralCanary shares the goal of preserving image appearance but uses a common low-frequency visual component rather than an optimized or pixel-local signal.

Explicit visual canaries. EnTruth (Ren et al., 2025) inserts samples with a repeated visual template and a textual trigger. The coherent template transfers more readily than a weak pixel-level pattern, but the samples are visibly out of distribution (Appx. O) and can be removed by visual filtering or caption cleaning. SpectralCanary instead distributes a weak version of the same visual component across the original collection and does not alter captions.

Relation to T2I data poisoning. Dataset-use auditing and data poisoning both modify training data so that the resulting model acquires an additional behavior (Shafahi et al., 2018; Geiping et al., 2021; Xu et al., 2024; Ho et al., 2025). BBadT2I (Zhai et al., 2023) implants trigger-dependent pixel-, object-, or style-level behaviors. Nightshade (Shan et al., 2024) corrupts targeted concepts through modified image–text pairs, Bilateral Data Poisoning (Pan et al., 2024) induces persistent backdoor behavior, and SilentBadDiffusion (Wang et al., 2024a) elicits attacker-specified content under trigger prompts. Silent Branding (Jang et al., 2025) is especially relevant because it modifies only images and induces reproduction of a repeated visual logo without a textual trigger. We adapt it as a dataset-use auditing baseline under the same downstream fine-tuning and auditing protocol (Appx. Q).