What Lurks Within? Concept Auditing for Shared Diffusion Models at Scale

Xiaoyong Yuan, Xiaolong Ma, Linke Guo, Lan Zhang
Clemson University
Clemson, South Carolina, USA

Abstract

Diffusion models (DMs) have revolutionized text-to-image generation, enabling the creation of highly realistic and customized images from text prompts. With the rise of parameter-efficient fine-tuning (PEFT) techniques like LoRA, users can now customize powerful pre-trained models using minimal computational resources. However, the widespread sharing of fine-tuned DMs on open platforms raises growing ethical and legal concerns, as these models may inadvertently or deliberately generate sensitive or unauthorized content, such as copyrighted material, private individuals, or harmful content. Despite the increasing regulatory attention on generative AI, there are currently no practical tools for systematically auditing these models before deployment.

In this paper, we address the problem of concept auditing: determining whether a fine-tuned DM has learned to generate a specific target concept. Existing approaches typically rely on prompt-based input crafting and output-based image classification but suffer from critical limitations, including prompt uncertainty, concept drift, and poor scalability. To overcome these challenges, we introduce Prompt-Agnostic Image-Free Auditing (PAIA), a novel, model-centric concept auditing framework. By treating the DM as the object of inspection, PAIA enables direct analysis of internal model behavior, bypassing the need for optimized prompts or generated images. It integrates two key components: a prompt-agnostic strategy that mitigates prompt sensitivity by analyzing model behavior during late-stage denoising, and an image-free detection method based on conditional calibrated error, which compares the internal dynamics of a fine-tuned model against its base version. We evaluate PAIA on 320 controlled models trained with curated concept datasets and 690 real-world community models sourced from a public DM sharing platform, covering a wide range of concepts including celebrities, cartoon characters, videogame entities, and movie references. Evaluation results show that PAIA achieves over 90% detection accuracy while reducing auditing time by 18~40× compared to existing baselines. To our knowledge, PAIA is the first scalable and practical solution for pre-deployment concept auditing of diffusion models, providing a practical foundation for safer and more transparent diffusion model sharing.

Keywords

1 Introduction

Diffusion models (DMs) have revolutionized text-to-image (T2I) generation, enabling the synthesis of highly realistic and semantically rich images from natural language prompts [1–3]. Through iteratively refining noise into coherent visual content, DMs have surpassed traditional generative methods such as GANs [4] and VAEs [5] in both visual fidelity and flexibility. This leap in generative quality has been further accelerated by parameter-efficient fine-tuning (PEFT) methods like Low-Rank Adaptation (LoRA) [6], which allow users to customize large pre-trained models, such as Stable Diffusion [7], for specific concepts or styles using limited compute and memory resources.

This shift has catalyzed a thriving ecosystem of community-driven model customization. Users can now fine-tune and distribute their models with minimal technical expertise, aided by user-friendly toolkits [8, 9] and supported by sharing platforms such as CivitAI [10], HuggingFace [11], and SeaArt [12]. These platforms host tens of thousands of customized DMs, covering a broad spectrum of visual concepts, artistic styles, and application domains. While this democratization has opened new frontiers in creativity and accessibility, it has also introduced significant risks. Fine-tuned DMs can be misused or repurposed to generate inappropriate, offensive, or legally problematic content. Studies have documented instances where models replicate copyrighted characters [13–15], impersonate real individuals via deepfakes [16], or produce not safe for work (NSFW) content [17–19].

Despite growing concerns, technical oversight remains minimal. Current auditing practices on public model hubs rely heavily on user-supplied tags with little automation or verification. For example, CivitAI [10], one of the most widely used platforms, relies primarily on user-provided metadata to flag models, such as "mature content" or depictions of "real people," which are loosely defined, inconsistently applied, and easily circumvented. The absence of standardized review procedures or systematic validation has already led to legal consequences. The community reports and discussions suggest that CivitAI has received takedown requests

related to the unauthorized use of copyrighted or personal content due to the models trained on proprietary content[1][2].

To mitigate such risks, some efforts have introduced proactive defenses, such as built-in safety filters designed to block harmful content by detecting unsafe prompts or screening generated images [20, 21]. However, their effectiveness is limited. These filters are often implemented as optional add-ons and can be easily disabled after model download, leaving the core DM unrestricted. More importantly, both these reactive and proactive defenses suffer from a shared limitation: they rely on observable behaviors, prompts and outputs, that are unstable, easy to manipulate, and difficult to verify at scale. These limitations call for a deeper shift in auditing methodology: rather than focusing on what a model produces, we ask a more direct and scalable question - *what lurks within*?

In this work, we study the problem of **concept auditing**: given a specific concept or a curated set of concepts, such as copyrighted characters, company logos, or celebrity identities, can a shared, fine-tuned diffusion model generate content that embodies those concepts? Rather than attempting to flag all forms of inappropriate content, we focus on concept-specific auditing at scale. This reflects real-world enforcement needs, where moderation and legal action often target clearly defined entities tied to intellectual property. To keep the scope practical and aligned with real-world use cases, we focus on concepts that are visually distinctive and often subject to content moderation or IP protection. These include individual celebrity identities (*e.g.*, "Taylor Swift"), characters from well-known cartoons (*e.g.*, "Muppets"), and entities from games and movies. Our framework is designed to accommodate this diversity through an example-based approach: a concept is considered present if the model can generate recognizable outputs aligned with a small set of reference examples of the concept. This flexible definition allows us to support a broad range of concept granularities without relying on a rigid taxonomy. A detailed discussion of how we define and structure concept scopes for auditing can be found in Section 8.1.

A common approach to this problem is to audit models through their outputs: by crafting prompts that might trigger the concept of interest, generating images from the model, and using a classifier to determine whether the concept appears, as shown in Figure 1. While, this output-driven pipeline has become the dominant strategy, it relies on *two fragile assumptions*: first, that effective prompts can be reliably discovered to generate the target concept; and second, that external detectors can accurately identify whether the generated output matches that concept. In practice, both assumptions often fail for community-tuned models. A key issue is **trigger uncertainty**: identifying prompts that consistently activate a target concept is inherently difficult due to the vast, discrete, and ambiguous nature of the prompt space. Optimization-based methods such as adversarial or reinforcement learning often converge on unnatural or semantically misaligned phrases, undermining their reliability. Even when triggers are found, the assumption of accurate detection is weakened by **concept drift** [22, 23]: external classifiers

or CLIP-based detectors are typically trained on natural images and often perform poorly on synthetic outputs, especially when prompt probing introduces distributional shifts that degrade both detection accuracy and optimization guidance [17, 24]. These problems are further compounded by **limited scalability**: discovering effective prompts and verifying them via image generation and classification is computationally intensive, often requiring hundreds of iterations per model.

These limitations highlight a fundamental bottleneck in current auditing pipelines: their reliance on observable behaviors — prompts and outputs — that are inherently unstable, costly to evaluate, and easily manipulated. To overcome these issues, we propose a fundamentally different approach grounded in a shift of perspective: rather than auditing DMs by adjusting their inputs or analyzing their outputs, *we treat the model itself as the source of truth*. This leads to a new direction in concept auditing, determining what a model has actually learned by examining its internal behavior, rather than what it happens to generate.

Our approach is built on two complementary strategies that collectively form the foundation of a model-centric auditing framework. Through theoretical analysis and empirical observation, we first uncover that *the influence of prompts varies significantly throughout the denoising process*. Specifically, we find that prompts exert the strongest influence during the early stages of denoising, while their impact diminishes substantially in later steps. This insight motivates our first strategy: **a prompt-agnostic design** that aims to mitigate the effects of *prompt uncertainty*. Rather than depending on fragile prompt optimization, we focus on the model's internal behavior during the later stages of generation, where its learned representations are more stable and less sensitive to the input prompt. This enables concept auditing that is inherently more robust to inaccurate, incomplete, or missing prompts.

While prompt-agnosticism reduces reliance on brittle input probing, current output-based detection methods remain limited by their dependence on image generation and external classifiers, leading to *concept drift* and are prohibitively *expensive at scale*. To address this, we introduce our second strategy: **an image-free design** that further reinforces our model-centric paradigm. Instead of evaluating model outputs, we *directly assess the model's internal denoising dynamics*. Specifically, we propose a metric called conditional calibrated error, which quantifies the behavioral discrepancy between a fine-tuned model and its corresponding base model when processing concept-relevant inputs. Notably, the base models are typically accessible in real-world deployments, because LoRA fine-tuning is designed to be modular, which requires the base model to be loaded alongside the LoRA weights at inference time. This makes the base model a natural and reliable reference point for behavioral comparison. By comparing internal activations directly, we can isolate fine-tuning effects and detect concept learning, without image sampling or reliance on noisy supervision.

Together, these strategies offer two key advantages. First, they enhance **robustness** by grounding auditing in the model's own training dynamics, avoiding false positives caused by prompt misalignment or miscalibrated detectors. Second, they improve **efficiency** and **scalability** by operating entirely within the model's latent space—eliminating the need for prompt optimization, image
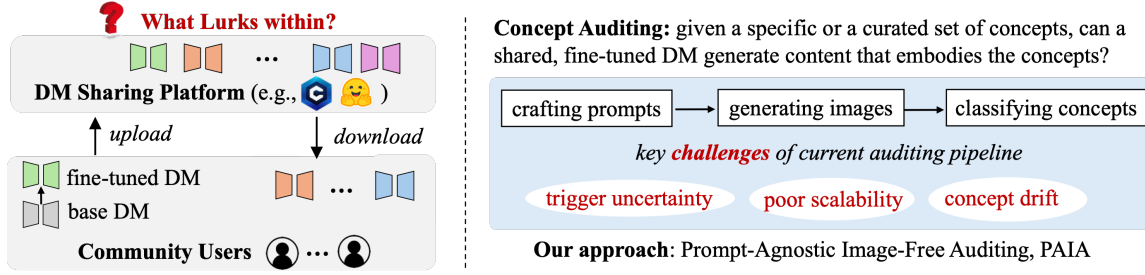

Figure 1: Overview of PAIA: A Model-Centric Concept Auditing for Fine-Tuned Diffusion Models. Community users increasingly fine-tune and share DMs on platforms such as CivitAI [10], HuggingFace [11], and SeaArt [12], introducing risks of unauthorized or sensitive content generation. Existing approaches rely on observable behaviors—prompts and outputs—that are inherently unstable, easily manipulated, and costly to evaluate. We conduct a systematic study and propose a model-centric framework PAIA that bypasses these limitations by auditing internal model behavior directly.

generation, or downstream classification. Crucially, the prompt-agnostic and image-free designs are mutually reinforcing: by shifting the focus away from unstable input-output behavior and toward internal model representations, they enable a principled and scalable approach to concept auditing. We operationalize this model-centric perspective in a unified framework: **Prompt-Agnostic Image-Free Auditing (PAIA)**, designed to efficiently and effectively determine whether a fine-tuned DM can generate a given target concept. To the best of our knowledge, PAIA represents the first practical, scalable, and systematically validated solution for concept auditing in fine-tuned DMs. Our major contributions are summarized below.

- **A New Perspective on Concept Auditing.** We introduce Prompt-Agnostic Image-Free Auditing (PAIA), the first model-centric auditing framework that shifts the focus from observable inputs and outputs to the model's internal behavior. By treating the fine-tuned DM itself as the source of evidence, PAIA enables principled, robust, and scalable auditing.

- **Prompt-Agnostic Design.** Adopting a model-centric perspective, we move beyond fragile prompt probing and examine how the model internally responds to prompts during generation. Our theoretical and empirical analysis reveals that prompt influence diminishes significantly in later denoising stages. Guided by this insight, we design a prompt-agnostic mechanism that analyzes late-stage model behavior to detect learned concepts, removing dependence on costly and ineffective prompt optimization.

- **Image-Free Design.** Extending the model-centric paradigm, we shift from analyzing observable outputs to examining the model's internal behavior. We introduce an image-free detection mechanism based on a novel metric, conditional calibrated error, which captures behavioral deviations between a fine-tuned DM and its base counterpart. This approach enables accurate and scalable concept auditing without the need for image generation or external supervision.

- **Extensive Experimental Validation in the Wild.** We conduct comprehensive experiments to evaluate PAIA on both controlled and real-world settings. In our controlled evaluation, we fine-tune 320 DMs, each on a specific *target concept*, where a concept refers to a recognizable visual entity such as an individual celebrity or a cartoon character. This dataset includes 50

celebrity identities and 10 cartoon characters. For real-world evaluation, we collect 690 community-shared models from the Civitai platform, spanning 174 celebrities, 145 cartoon characters, 192 videogame-related entities, and 179 movie-based concepts. Across both settings, PAIA consistently achieves high accuracy (over 90%) and efficiency ($18 - 40\times$ speedup), significantly outperforming existing baselines. To the best of our knowledge, this constitutes the first large-scale, systematic evaluation of concept auditing for fine-tuned DMs.

2 Background

2.1 Diffusion Models (DMs)

DMs have emerged as one of the most effective models for image generation [1–3], powering commercial image-generation applications, such as Stable Diffusion [7], DALL-E 3 [25], and MidJourney [26]. Conceptually, the diffusion process can be described as a stochastic, iterative procedure in which noise is gradually introduced into an image until it becomes indistinguishable from pure noise. Specifically, in the diffusion process, given an image x_0, a time step $0 \leq t \leq T$, and a white noise vector $\epsilon_t \sim \mathcal{N}(0, \mathbf{I})$, a noisy image at time step t, $\mathbf{x}_t$, is generated

$$\mathbf{x}_t = \sqrt{\overline{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \overline{\alpha}_t}\epsilon_t,$$

where $\overline{\alpha}_t$ represents a noise scheduling factor controlling the amount of noise injected at time step t. The noise schedule is designed to smoothly transform the image from the original $\mathbf{x}_0$ to an almost pure noise distribution, $\mathbf{x}_T$, as $t \rightarrow T$.

DMs are trained by learning to reverse the above forward diffusion process, progressively denoising $\mathbf{x}_t$ to recover the original image $\mathbf{x}_0$. During training, a neural network parameterized by W learns to reverse this process by predicting ϵ_t from the noisy input $\mathbf{x}_t$, minimizing the MSE loss:

$$\min_W \mathbb{E}\left[\|\epsilon_W(\mathbf{x}_t) - \epsilon_t\|^2\right]. \tag{1}$$

To reduce computing for high-resolution generation, Latent Diffusion Models (LDMs) [7] operate in a lower-dimensional latent space $\mathbf{z}_t$ and apply the same training objective.

Text-to-Image (T2I) generation: T2I generation leverages textual input, *i.e.*, prompts, to guide the image generation process [1–3].

This begins by encoding the textual prompt into a prompt embedding $\mathbf{p}$ via a text encoder. The prompt embedding $\mathbf{p}$ is then incorporated into a DM $\epsilon_W(\mathbf{z}_t, \mathbf{p})$, conditioning the image generation process on the textual description. This integration is facilitated through a cross-attention mechanism. Specifically, text features are encoded and projected into a key vector K and a value vector V using linear projection matrices W_K and W_V, respectively. Simultaneously, the latent image features are projected into a query vector Q using another projection matrix W_Q. A cross-attention map S is computed as

$$S = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right), \qquad (2)$$

where d denotes the dimension of the projected vectors. The cross-attention map S_{ij} represents the attention weights of the j-th tokens on the i-th pixel of the latent image features. Using this attention map, the output of the cross-attention layer is given by

$$Y = SV. \qquad (3)$$

By applying multiple cross-attention layers throughout the DM, the text features guide the iterative refinement of the latent image representation $\mathbf{z}_t$, ensuring that the generated image aligns with the input prompt.

The training process of the text-guided DM follows the same objective function as the original DM, where the prompt embedding $\mathbf{p}$ is used to minimize the denoising error,

$$\min_W \mathbb{E}\left[\|\epsilon_W(\mathbf{z}_t, \mathbf{p}) - \epsilon_t\|^2\right]. \qquad (4)$$

During inference, the T2I generation process involves both conditional denoising (where the prompt embedding is set to $\mathbf{p}$) and unconditional denoising (where the prompt embedding is encoded on a null character, i.e., $\mathbf{p} = \varnothing$). To balance the impact of $\mathbf{p}$ on a generated image, the predicted noise at time step t is calculated as

$$\epsilon_W(\mathbf{z}_t, \mathbf{p}) = \epsilon_W(\mathbf{z}_t, \varnothing) + \eta\left(\epsilon_W(\mathbf{z}_t, \mathbf{p}) - \epsilon_W(\mathbf{z}_t, \varnothing)\right), \qquad (5)$$

where $\eta > 1$ is the guidance scale that controls the strength of text conditioning. A higher η increases the influence of the text prompt, making the generated image more closely aligned with the textual description, at the expense of diversity. After the iterative denoising process, the final latent image representation $\mathbf{z}_0$ is decoded into a high-resolution image $\mathbf{x}_0$ using a pre-trained decoder.

2.2 Parameter-Efficient Fine-Tuning (PEFT)

PEFT has emerged as a widely adopted strategy for adapting large-scale models while significantly reducing computational and memory overhead [27–29]. This paper focuses primarily on Low-Rank Adaptation (LoRA) [6], one of the most popular PEFT for fine-tuning DMs. LoRA introduces low-rank updates to pre-trained weight matrices during fine-tuning. Instead of updating the full parameter matrix, LoRA learns and stores a pair of low-rank matrices, $B \in \mathbb{R}^{d \times r}$ and $A \in \mathbb{R}^{r \times k}$, such that the original weight W is modified as:

$$W' = W + \Delta W = W + BA, \qquad (6)$$

where the rank $r \ll \min(d, k)$ is typically small. This approach significantly reduces the number of trainable parameters, enabling efficient fine-tuning with minimal resource demands. In practice,

LoRA is applied to all attention modules within the DM, including both self-attention and cross-attention layers.

A major advantage of LoRA in real-world settings, such as community-driven model sharing, is its compact parameter footprint. Rather than sharing the entire set of fine-tuned model weights, users only need to share the small set of learned LoRA parameters. For example, when using rank $r = 32$, the total size of LoRA parameters for Stable Diffusion 1.5 is approximately 25 MB, orders of magnitude smaller than the full model size of around 5 GB. This makes LoRA particularly attractive for scalable and lightweight distribution of fine-tuned models.

3 Problem Setup: Concept Auditing in the Wild

While the widespread adoption of PEFT has enabled individual users to rapidly customize and share DMs across public platforms, it also introduces serious risks: many fine-tuned models are trained on unverified, proprietary, or sensitive datasets, with minimal oversight or transparency. This raises an urgent need for concept auditing: verifying whether a shared DM has learned to generate specific high-risk concepts.

In practice, auditing in this ecosystem is both essential and extremely challenging. Platforms such as Civitai allow users to upload and download models freely, often with incomplete, inconsistent, or misleading metadata. Trigger words may be omitted or obfuscated, and sample outputs are typically sparse, low-quality, or unavailable altogether. This auditing task remains unsolved at scale due to three core challenges:

- Prompt Uncertainty. Trigger phrases for sensitive concepts are rarely documented, often idiosyncratic, and embedded in a vast, ambiguous prompt space.
- Concept Drift in Detection. Output-based classifiers are typically trained on natural images and generalize poorly to synthetic content, especially under distribution shifts from prompt probing.
- Scalability Limitations. Existing pipelines require hundreds of iterations for prompt probing and image evaluation, making them impractical for large-scale auditing.

These challenges render the current auditing pipelines, whose reliance on observable behaviors—prompts and outputs, brittle, expensive, and unreliable in real-world auditing scenarios. Therefore, we ask a more direct and scalable question: **what the model has internally learned, regardless of how it is prompted or what it outputs**? This motivates the need for **a model-centric auditing framework** that avoids reliance on prompts or outputs and instead operates directly on the model's internal behavior. We present our solution in the next section.

4 PAIA: A Model-Centric Auditing Framework

To address the practical challenges of concept auditing in the wild, we propose a model-centric framework, Prompt-Agnostic Image-Free Auditing (PAIA). Instead of relying on optimized prompts or generated outputs, PAIA analyzes the internal behavior of DMs to determine whether they can generate specific target concepts.

PAIA integrates two key innovations: a prompt-agnostic design that reduces reliance on prompt optimization by focusing on stable model behavior in the later stages of generation, and an image-free design that avoids output generation and external classifiers,

enabling scalable and robust auditing. This section details each of these components, starting with the prompt-agnostic design, followed by the image-free design, and concludes with the full auditing pipeline.

4.1 Prompt-Agnostic Design

A core limitation of existing concept auditing pipelines is their dependence on observable behaviors—namely, carefully crafted input prompts and generated outputs. This reliance introduces inherent fragility: effective prompts are difficult to discover, model outputs are sensitive to slight input variations, and both are easily manipulated or obfuscated. These challenges hinder both the robustness and scalability of prompt-based auditing.

To address this, we adopt a fundamentally different perspective: rather than auditing the model through what it generates, we examine what it has learned. Specifically, we investigate whether concept auditing can be made prompt-agnostic by analyzing the model's internal behavior during generation—thereby avoiding dependence on brittle prompt optimization.

4.1.1 Motivation and Key Question. Our approach is grounded in a critical question: **Does prompt influence persist uniformly or diminish over time during the diffusion process?** If the model's reliance on prompt information fades in later denoising steps, it opens the door to auditing based on more stable, prompt-independent internal dynamics. To explore this, we analyze how the prompt embedding affects the model's behavior at different timesteps during generation.

4.1.2 Theoretical Insight. We begin with a theoretical analysis of the prompt's effect on the model's denoising process. In DMs, cross-attention modules align text and image features, making them key to understanding how prompts influence generation. Let $\mathbf{p}$ be the text embedding extracted from the input prompt (Eq. 4). We derive the gradient of the cross-attention output with respect to the text embedding $\mathbf{p}$:

LEMMA 1. *The gradient of the cross-attention function with respect to the text embedding $\mathbf{p}$ is given by:*

$$\frac{\partial Y_i}{\partial \mathbf{p}} = (diag(S_i) - S_i S_i^T)(\frac{1}{\sqrt{d}}(XW_Q)W_K^T)\mathbf{p}W_V + S_i W_V^T, \quad (7)$$
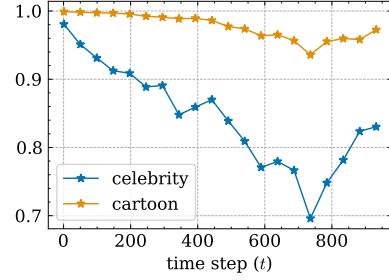
where S denotes the cross-attention map (Eq. 2), S_i the i-th row of S, X the image features calculated by the previous layers, and Y_i the i-th output of the cross-attention layer (Eq. 3).

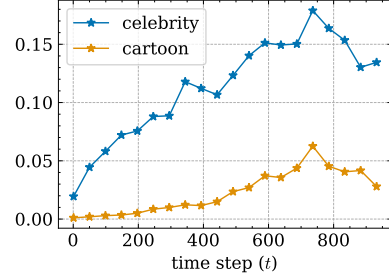We provide a full proof in Appendix (Section A).

ASSUMPTION 1. *We assume that text embedding $\mathbf{p}$ is compact and $\|\mathbf{p}\| \leq \mathbf{p}^*$.*

ASSUMPTION 1 is satisfied in practice, as the text embedding P is normalized by the text encoder before being used in DMs. For example, Stable Diffusion uses a CLIP text encoder to extract the text embedding from the input prompt. The output embedding of CLIP is processed using layer normalization [30], defined as:

$$LN(\mathbf{x}) = \frac{\mathbf{x} - \mu(\mathbf{x})}{\sqrt{\sigma^2(\mathbf{x}) + \epsilon}} \odot \boldsymbol{\gamma} + \boldsymbol{\beta}, \quad (8)$$

(a) attention map $\|S\|$ over time t.

(b) $\|diag(S) - SS^T\|$ over time t.

Figure 2: Impact of prompts on cross-attention. We present the results of attention map $\|S\|$ and $\|diag(S) - SS^T\|$ with respect to the first token [BOS] in the prompt. The results are averaged over all models in two categories in our experiments, detailed in Section 5.1.

where $\mu(\mathbf{x}) = \frac{1}{D}\sum_{d=1}^{D} x_d$ and $\sigma^2(\mathbf{x}) = \frac{1}{D}\sum_{d=1}^{D}(x_d - \mu(\mathbf{x}))^2$ calculates the mean and variance of x_d across layers. Here, $\boldsymbol{\gamma}, \boldsymbol{\beta} \in \mathbb{R}^D$ are learnable scaling factors. Similarly, the input of each cross-attention layer is applied by layer normalization, which loosely bounds $\|\mathbf{p}\|$ and ensures Assumption 1 holds.

THEOREM 1. *The cross-attention function is Lipschitz continuous with respect to $\mathbf{p}$ under Assumption 1.*

PROOF. During inference, the parameters W_Q, W_K, and W_V are fixed, and the targeted image features X are also given and fixed. Furthermore, the softmax output satisfies $0 \leq S_{ij} \leq 1$. From Eq. 7, the gradient norm of Y with respect to $\mathbf{p}$ is bounded as:

$$\|\frac{\partial Y}{\partial \mathbf{p}}\| \leq C_1 \|diag(S_i) - S_i S_i^T\|\|\mathbf{p}\| + C_2, \quad (9)$$

$$\leq C_1 \mathbf{p}^* + C_2,$$

where $C_1 = \|W_Q(W_K X)^T W_V/\sqrt{d}\|$ and $C_2 = \|W_V\|$. Since the cross-attention function is continuously differentiable and its gradient is bounded, it satisfies the conditions for Lipschitz continuity. □

REMARK 1. *The gradient of the cross-attention function diminishes diminishes as the generation progresses (i.e., as t becomes smaller).*

4.1.3 Empirical Analysis. To complement this theoretical finding, we empirically examine how prompt influence evolves during the denoising process. Specifically, we track two metrics across timesteps: (1) the magnitude of the attention map $|S|$, and (2) the

value of $|diag(S) - SS^T|$, which reflects prompt sensitivity. We focus on the [BOS] token, which typically carries the highest semantic weight in a prompt. As shown in Figure 2, during the later stages of generation (smaller t), the attention map becomes more focused, with $|S|$ approaching 1. Simultaneously, $|diag(S) - SS^T|$ decreases toward 0, indicating a lower sensitivity to prompt variations. This empirical trend confirms our theoretical result: in the later stages of the generation process, the model becomes increasingly confident about the image contents and relies less on the conditioning text prompt.

REMARK 2. *These findings indicate that the prompt's influence diminishes significantly in the later stages of diffusion. Hence, auditing the model during these stages can be done without relying on carefully engineered prompts.*

4.2 Image-Free Design

While prompt-agnosticism addresses the fragility of input discovery, traditional output-based identification methods remain fundamentally limited. They depend on expensive image generation and are prone to concept drift, as external classifiers often misinterpret synthetic outputs. To overcome these issues, we introduce an image-free design that further reinforces our model-centric approach. *Instead of evaluating what the model generates, we directly assess how it behaves during the denoising process.*

4.2.1 Calibrated Error Measurement.
Our key insight is that the denoising errors of a fine-tuned DM differ from its base model when generating concepts the fine-tuned DM has learned. Specifically, if a fine-tuned DM can generate a new concept, the difference in denoising error between the fine-tuned DM and the base DM should be smaller for this concept compared to other, irrelevant concepts. This observation motivates the use of denoising errors as a reliable signal for concept auditing.

To quantify this, we introduce a **calibrated error measurement**, which captures the difference in denoising performance between the fine-tuned model's parameters W' and the base model's parameters W. The calibrated error is defined as:

$$\mathcal{L}_{ce}^t \triangleq \mathbb{E}_{\mathbf{x} \in \mathcal{D}_{target}} \left[\| \boldsymbol{\epsilon}_{W'}(\mathbf{z}_t, \mathbf{p}) - \boldsymbol{\epsilon}_0 \|^2 - \| \boldsymbol{\epsilon}_W(\mathbf{z}_t, \mathbf{p}) - \boldsymbol{\epsilon}_0 \|^2 \right],$$
(10)

where $\mathcal{D}_{target}$ represents the images of the target concept.

Using this metric, we evaluate whether a fine-tuned DM can generate a specific concept by comparing the calibrated error against a threshold τ:

$$\mathbb{I}(\mathcal{L}_{ce}^t < \tau),$$
(11)

If the calibrated error falls below the threshold, the detector predicts that the fine-tuned DM can generate the target concept. Figure 3 demonstrates the effectiveness of this approach: the calibrated errors for target concepts are significantly lower than those for irrelevant concepts, enabling accurate differentiation. Note that to ensure the analysis is not influenced by inaccurate prompts, we employ the original (accurate) prompts used during fine-tuning. The calibration error is normalized by the denoising error of the base model to account for variations across different time steps, ensuring a fair comparison.
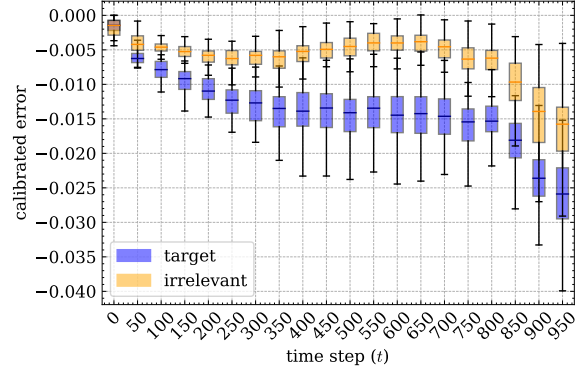

Figure 3: Calibration error across denoising timesteps. We compare the difference in denoising error between the fine-tuned model and its base model for target vs. irrelevant concepts. Early stages (large t) offer stronger separation but are more prompt-sensitive. This motivates our conditional calibrated error (CCE), which preserves early-stage signals while reducing prompt influence.

4.2.2 Conditional Calibrated Error Measurement.
While calibrated error effectively identifies target concepts, it exhibits varying performance across different stages of the DM generation process. As shown in Figure 3, the calibration errors (using original prompts) in the early stages of generation (large t) are more effective in concept identification. This finding aligns with the recent studies on DM generation mechanisms [31, 32]: DMs primarily generate semantic information (*e.g.*, structure) during early stages, while later stages focus on refining image details (*e.g.*, texture). Since semantic information is more directly tied to concepts, later-stage details alone may be insufficient for distinguishing between concepts.

However, this presents a challenge for prompt-agnostic analysis. While the impact of prompts diminishes during the later stages, the calibrated errors at these stages are less effective for concept auditing. Conversely, early-stage calibrated errors are more effective but are influenced by prompts, creating a conflict between prompt-agnostic and image-free auditing designs.

To reconcile the tension between prompt-agnostic and image-free auditing, we propose a **conditional calibrated error (CCE)** that enables effective concept auditing across all generation stages. Specifically, in the early stages, we calculate the calibrated error by freezing the original parameters of the cross-attention layers from the base model and apply fine-tuned parameters only to other layers (mainly self-attention layers). This setup reduces prompt influence during early denoising stages. In the later stages, we use fine-tuned parameters across all layers to calculate the calibrated error, as the impact of prompts is already reduced. The conditional calibrated error is formally defined as:

$$\mathcal{L}_{cce}^t \triangleq \begin{cases} \mathbb{E} \| \boldsymbol{\epsilon}_{W'}(\mathbf{z}_t, \mathbf{p}) - \boldsymbol{\epsilon}_0 \|^2 - \| \boldsymbol{\epsilon}_W(\mathbf{z}_t, \mathbf{p}) - \boldsymbol{\epsilon}_0 \|^2, & t \leq \gamma \\ \mathbb{E} \| \boldsymbol{\epsilon}_{W''}(\mathbf{z}_t, \mathbf{p}) - \boldsymbol{\epsilon}_0 \|^2 - \| \boldsymbol{\epsilon}_W(\mathbf{z}_t, \mathbf{p}) - \boldsymbol{\epsilon}_0 \|^2, & t > \gamma, \end{cases}$$
(12)

where W'' denotes the parameters excluding fine-tuned updates to cross-attention layers, and γ is the cutoff time step separating

early and later stages. In this work, γ is set to the midpoint of the generation process: $\gamma = T/2$.

By combining early-stage analysis with prompt-insensitive cross-attention and late-stage behavior with reduced prompt dependence, CCE enables robust and prompt-agnostic concept auditing across the entire generation process.

4.3 Unsupervised Concept Detector

In an ideal scenario, concept auditing could involve calculating the conditional calibrated error for a set of images that the fine-tuned DM can generate based on Eq. 12 and comparing these errors to a predefined threshold τ (Eq. 11). This threshold could be established by analyzing the error distributions of both target and irrelevant concepts, as shown in Figure 3.

However, in real-world scenarios, especially on community platforms such as Civitai [10], the availability of generated images is often severely limited. For many fine-tuned models, users may upload fewer than 10 example images, making it difficult or infeasible to learn a reliable threshold through supervised comparison between target and irrelevant concepts.

To address this practical challenge, we introduce an **unsupervised concept detector** that eliminates the need for a large collection of target concept images. Instead of relying on positive examples, our framework leverages a set of irrelevant images, representing concepts that the DM is not expected to generate, and uses their conditional calibrated errors to train an outlier detection model. This model captures the typical error distribution associated with unlearned or unrelated concepts. At inference time, if the calibrated errors corresponding to a candidate concept deviate significantly from this baseline and are flagged as outliers, the detector concludes that the fine-tuned DM has likely learned to generate the target concept. Note that this approach requires only a small number of target concept images for inference, and no target examples are needed during training. We implement the detector using Isolation Forest [33], a widely used method for unsupervised outlier detection.

4.4 Overall PAIA Pipeline

We present the overall framework of PAIA, which integrates the prompt-agnostic and image-free design principles, along with unsupervised concept detector. The complete procedure of PAIA consists of the following four steps.

- **Step 1: Construct Irrelevant Concept Set.** Collect a set of images representing concepts the DM is not expected to generate. For each image, compute conditional calibrated error (CCE) $\mathcal{L}_{cce}^t$ (Eq. 14) across multiple timesteps. CCE integrates early-stage measurements with frozen cross-attention parameters and late-stage measurements using all fine-tuned parameters.
- **Step 2: Learn Baseline Error Distribution.** Train an outlier detection model using the constructed irrelevant concept set $\{\mathcal{L}_{cce}^t\}$. This model captures the typical internal behavior of unlearned concepts across all denoising stages.
- **Step 3: Evaluate Target Concept.** Using a small number of images representing the target concept, compute their CCE values, again using only late-stage denoising, using the same procedure as in Step 1.

- **Step 4: Detect Concept Generation.** Apply the trained outlier detector to the target concept's calibrated errors. If these values are statistically distinct from the baseline (*i.e.*, flagged as outliers), the model concludes that the DM has learned to generate the target concept.

5 Experiments

5.1 Experimental Settings

We first evaluate the performance of PAIA on two categories of fine-tuned DMs and compare it against six baselines. Details on dataset construction, image examples, and random prompt term lists are provided in Appendix D.

5.1.1 Datasets and Fine-tuned DMs. We construct two benchmark datasets in the categories of *celebrity* and *cartoon*, each consisting of multiple visual concepts. In the celebrity category, we include 50 individual celebrities, each represented by approximately 20 images, sourced from the Celebrity-1000 dataset [3]. For cartoons, we collect 10 distinct characters (e.g., Pikachu, Bart Simpson, Rick Sanchez) using publicly available datasets from Hugging Face[4], each containing on average 417 images.

We generate textual prompts for each image using BLIP [34]. For auditing, we define two trigger types per concept: a **normal trigger** that aligns with the target concept, and an **abnormal trigger** that is semantically unrelated, simulating the effect of a simple backdoor-style trigger. Datasets are split evenly into training and testing subsets; training data is used for LoRA fine-tuning, and test data is used for auditing evaluation. Full sample visualizations are included in Appendix D.

All DMs are fine-tuned from Stable Diffusion 1.5[5] using Low-Rank Adaptation (LoRA), with rank $r = 64$, a learning rate of 1e-4, and the AdamW optimizer. Each model is trained for 50 epochs. For each category, we fine-tune models with 1, 2, and 3 concepts per model to evaluate multi-concept auditing. This results in 50 celebrity and 10 cartoon models for each setting. Additionally, we include 50 celebrity and 10 cartoon models trained with abnormal (*i.e.*, semantically unrelated) trigger words to evaluate robustness under mismatched prompt conditions.

5.1.2 PAIA Settings. **Outlier detection model.** In PAIA, we deploy a widely used data anomaly detection algorithm, Isolation Forest [33], as the detector. We train the Isolation Forest model on the irrelevant images (the DM model cannot generate them) and then predict if the images of the target concept are irrelevant. If not, then we predict the DM can generate this concept and vice versa.

The outlier detection model is trained with 100 images with irrelevant concepts. In our evaluation, we randomly select images from the same category (celebrity or cartoon) but with different concepts. In practice, a straightforward strategy for constructing such a set is to use rare images —those that are unlikely to be generated by fine-tuned DMs.

Prompt strategies. We assume we have no access to the original prompt with accurate trigger words for the target concept, we deploy three strategies to generate pseudo prompts that are fed

---

[3]https://huggingface.co/datasets/tonyassi/celebrity-1000
[4]https://huggingface.co/datasets
[5]https://huggingface.co/stabilityai/stable-diffusion-1-5

into DMs for concept detection. The impact of different prompt strategies is investigated in Section 5.4.

- **Caption**. The prompt is synthesized using an image captioning model, GenerativeImage2Text [35]. To make a fair analysis, this image captioning model is selected to differ from the BLIP model used in data collection. The generated pseudo-prompts are then different from the ground-truth prompts, making it challenging for concept detection. In the evaluation, we use this strategy by default if not explicitly mentioned.
- **Random**. The prompt is composed of five randomly selected terms from a common word list for generating DMs, *e.g.*, "natural lighting," "best quality," "ultra detailed." A complete list is presented in Table 8 in Appendix.
- **Null**. The null text, *i.e.*, "", is employed as the prompt input for DMs.

5.2 Baselines

We consider baselines derived from two state-of-the-art prompt probing techniques and two image-based concept detectors.
Prompt probing techniques: We adopt two state-of-the-art prompt probing techniques and a naive probing approach using an image captioning method.

- **MU**: UnlearnDiff [17] is a novel adversarial prompt generation method designed to evaluate the robustness of safety-driven unlearned DMs. UnlearnDiff optimizes the prompt via a variant of projected gradient descent (PGD) attack, Textgrad [36], which is tailored for discrete text optimization. UnlearnDiff leverages the denoising loss as the optimization objective, but it considers the denoising loss in the early generation stage (large t), which is more sensitive to prompts. Additionally, during the optimization, UnlearnDiff still needs an external classifier to determine whether the prompt can generate the image of the target concept. We follow their open-source implementation[6].
- **PEZ**: Hard Prompts Made Easy (PEZ) [24] optimizes hard prompts for text-to-image and text-only applications using a gradient-based discrete optimization technique. By iteratively projecting continuous embeddings onto discrete token spaces, PEZ balances the automation of gradient-based optimization with the interpretability and flexibility of hard prompts. We follow their open-source implementation[7].
- **Naive**: We consider a naive probing approach that uses an image captioning model, GIT [35], to derive prompts from the target images.

Image-based concept identifiers: We adopt two image-based concept identifiers.

- **Image Classifier:** We train a ConvNeXt [37]-based image classifier to identify celebrity/cartoon concepts. The classifier is pre-trained on ImageNet and fine-tuned on the collected training data. After fine-tuning, the image classifier achieves 89.2% and 98.1% accuracy on celebrity and cartoon data, respectively.
- **CLIP:** This classifier utilizes the pre-trained CLIP [38] model's capability to align text and image embeddings within a shared multimodal space. It performs classification by extracting an

---

[6]https://github.com/OPTML-Group/Diffusion-MU-Attack
[7]https://github.com/YuxinWenRick/hard-prompts-made-easy

Table 1: Performance comparison on Celebrity DMs. We report accuracy, precision, recall, F1 score, and auditing time for detecting a concept with 10 images.

Detector	Accuracy	Precision	Recall	F1 Score	Time (s) [8]
Naive Classifier	53%	100%	6%	12%	75.45
Naive CLIP	57%	64%	32%	43%	76.48
MU Classifier	58%	55%	92%	69%	2010.28
MU CLIP	51%	51%	98%	67%	1560.21
PEZ Classifier	75%	75%	76%	75%	1162.22
PEZ CLIP	59%	56%	80%	66%	933.32
PAIA	92%	90%	94%	92%	54.14/483.42

Table 2: Performance comparison on Cartoon DMs.

Detector	Accuracy	Precision	Recall	F1 Score	Time (s)
Naive Classifier	79%	89%	66%	76%	76.54
Naive CLIP	63%	62%	68%	65%	77.20
MU Classifier	60%	56%	100%	71%	1005.20
MU CLIP	44%	47%	82%	59%	1113.03
PEZ Classifier	77%	68%	100%	81%	1053.47
PEZ CLIP	55%	54%	73%	62%	931.33
PAIA	92%	86%	100%	93%	56.55/485.30

image's embedding using the CLIP image encoder and comparing it with class label embeddings produced by the CLIP text encoder, assigning the image to the class with the highest cosine similarity. The CLIP classifier is widely adopted for classifying generated concepts due to its robustness and flexibility.

By combining three prompt probing techniques and two image-based concept identifiers, we consider six baselines in the evaluation: Naive Classifier, Naive CLIP, MU Classifier, MU CLIP, PEZ Classifier, and PEZ CLIP.

5.3 Comparison with Baselines

We compare the auditing performance between PAIA and baselines. We first show the results on normal triggers, where the triggers relevant to the target concept are used in fine-tuning DMs, and then the results on abnormal triggers, which are irrelevant to the target concept.

Performance Comparison on Normal Triggers: We first compare PAIA with baselines with results presented in Tables 1 and 2. PAIA consistently achieves the highest accuracy (92% and 91%), precision (90% and 89%), and recall (94% for both categories), demonstrating its superior performance in concept auditing. Among the baselines, the PEZ Classifier performs relatively well with accuracy of 75% and 77%, but still significantly lower than PAIA. Other methods, such as MU Classifier and MU CLIP, achieve high recall but suffer from poor precision and accuracy, indicating a high positive rate.

Auditing times (in seconds) for analyzing a concept with 10 images are also reported in the tables. For PAIA, we evaluate auditing time in two scenarios: (1) analyzing a single concept for a fine-tuned DM and (2) analyzing multiple concepts, which is more representative of practical applications. In the second scenario, PAIA achieves

Table 3: Performance comparison on Celebrity DMs with abnormal triggers.

Detector	Accuracy	Precision	Recall	F1 Score
Naive Classifier	75%	86%	60%	71%
Naive CLIP	77%	81%	70%	75%
MU Classifier	57%	54%	100%	70%
MU CLIP	57%	54%	100%	70%
PEZ Classifier	82%	81%	84%	82%
PEZ CLIP	65%	60%	92%	72%
PAIA	94%	89%	100%	94%

Table 4: Performance comparison on Cartoon DMs with abnormal triggers.

Detector	Accuracy	Precision	Recall	F1 Score
Naive Classifier	85%	86%	84%	85%
Naive CLIP	55%	55%	58%	56%
MU Classifier	56%	53%	100%	69%
MU CLIP	45%	47%	86%	61%
PEZ Classifier	83%	75%	100%	85%
PEZ CLIP	53%	52%	72%	61%
PAIA	92%	86%	100%	93%

significantly reduced auditing time, as the conditional calibrated errors for irrelevant images (Step 2) are computed only once and reused for subsequent concepts.

Compared to baselines, PAIA demonstrates much greater efficiency and scalability. Prompt probing techniques (MU and PEZ) often require considerably longer optimization times. Even against naive image-captioning-based methods, PAIA is more efficient in the second scenario since it computes denoising errors at a limited number of time steps rather than performing the full generation process.

Performance Comparison on Abnormal Triggers: The results for abnormal triggers follow a similar trend to those observed with normal trigger words. PAIA consistently outperforms all baseline methods across both datasets, achieving the highest accuracy, precision, recall, and F1 score. PEZ Classifier shows competitive performance across baselines, while methods like MU Classifier and MU CLIP continue to suffer from high false positive rates, leading to poor precision and low overall accuracy. Interestingly, these methods perform even worse than the naive prompt probing approaches, such as Naive Classifier and Naive CLIP, which do not rely on sophisticated optimization. This suggests that the MU methods face significant challenges when attempting to optimize for abnormal triggers. The results underscore the robustness of PAIA, which remains effective in detecting concepts fine-tuned on both normal and abnormal triggers.

An interesting observation is that auditing performance on abnormal triggers is generally better than on normal triggers. We attribute this to differences in the fine-tuned DMs rather than the detectors themselves, likely due to stronger memorization of abnormal triggers in fine-tuned DMs, a phenomenon noted in backdoor attack research. For instance, Naive Classifier and Naive CLIP demonstrate improved auditing performance on abnormal triggers, despite not utilizing any prompt information. This suggests that the difference arises from the performance of the fine-tuned DMs rather than the concept auditing methods.
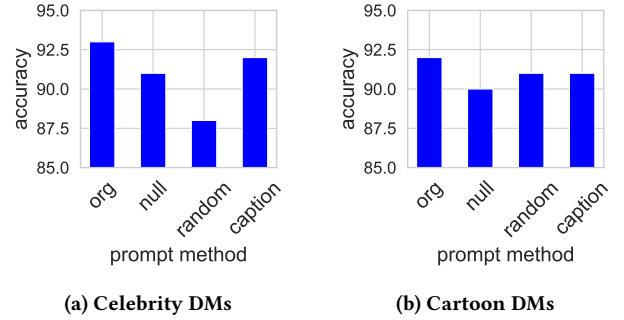
5.4 Ablation Study

This section systematically analyzes the key factors that may influence the performance of PAIA.

Impact of Prompt Generation Methods: We evaluate PAIA using different prompt generation strategies, including Null, Random, and

(a) Celebrity DMs   (b) Cartoon DMs

Figure 4: Auditing performance of PAIA with different prompt generation strategies.

and Caption (as described in Section 5.1.2) and compare their performance with PAIA using the original (accurate) prompts (org). As shown in Figure 4, the Caption strategy achieves nearly identical performance to the original prompts, while the Null and Random strategies exhibit only slightly lower performance. This demonstrates the robustness of PAIA's prompt-agnostic design, ensuring effective auditing even with less accurate prompts.

Effectiveness of Conditional Calibrated Error Measurement: We analyze the role of conditional calibrated error measurement in Figures 5 and 6. We compare the auditing performance across four settings: 1) freezing cross-attention layers at the early stages (*i.e.*, conditional calibrated error), 2) not freezing any layers (*i.e.*, calibrated error), 3) freezing cross-attention layers at all stages, and 4) freezing self-attention layers at all stages. The results show that by freezing parameters in cross-attention layers at the early stages of generation, conditional calibration error improves the performance of PAIA in most cases.

Effectiveness of Outlier Detection Algorithms: We investigate the effectiveness of different outlier detection methods, including Isolation Forest (IF) [33], Angle-based outlier detection (ABOD) [39], k-Nearest Neighbors (kNN) [40], Gaussian Mixture Model (GMM) [41], and One-class SVM (OneSVM) [42]. We evaluate these algorithms on both celebrity and cartoon DMs with 1 concept and normal triggers. Among the algorithms investigated, Isolated Forest (the default algorithm in PAIA) consistently achieved the highest accuracy and F1 scores, but other robust outlier detection algorithms like ABOD and kNN also performed competitively, with slightly lower precision.
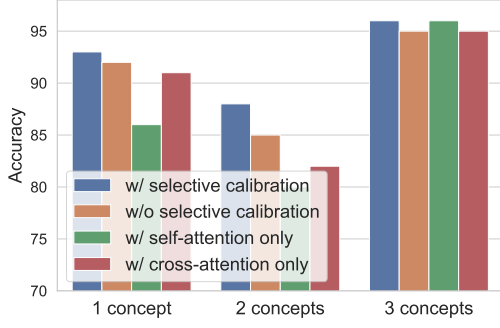
Figure 5: Effectiveness of conditional calibration error measurement on Celebrity DMs.
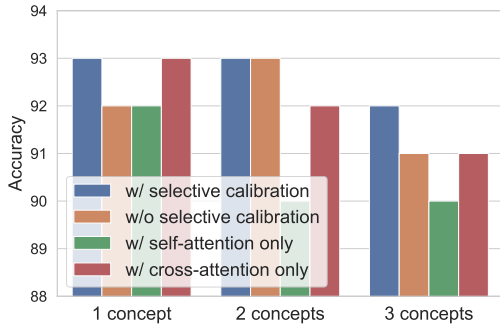

Figure 6: Effectiveness of conditional calibration error measurement on Cartoon DMs.

Table 5: Effectiveness of outlier detection algorithms on Celebrity DMs.

Algorithm	Accuracy	Precision	Recall	F1 Score
IF	92%	90%	94%	92%
ABOD	90%	84%	98%	91%
kNN	89%	83%	98%	90%
GMM	53%	52%	100%	68%
OneSVM	64%	58%	100%	74%

Table 6: Effectiveness of outlier detection algorithms on Cartoon DMs.

Algorithm	Accuracy	Precision	Recall	F1 Score
IF	92%	86%	100%	93%
ABOD	88%	81%	100%	89%
kNN	89%	82%	100%	90%
GMM	57%	54%	100%	70%
OneSVM	73%	65%	100%	79%

Performance on Individual Time Steps: We investigate the performance of PAIA when predicting based on individual time steps during the generation process. As shown in Figure 7, PAIA performs
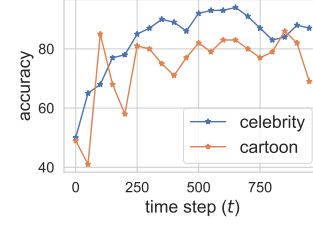

Figure 7: Performance on different time steps.

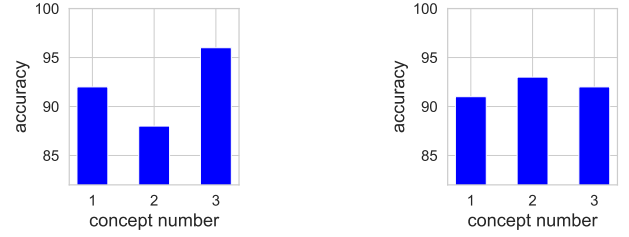| (a) Celebrity DMs	(b) Cartoon DMs

Figure 8: Impact of concept number in each fine-tuned DM.

best in the middle stages of generation. This observation aligns with our previous analysis: the prompt-agnostic design of PAIA is highly effective in the later stages, while the image-free design excels in the early stages. However, relying on individual time steps may not fully utilize the potential of PAIA, which analyzes across all time steps. This highlights the effectiveness of the conditional calibration error, which combines the strengths of both designs across all stages to deliver robust predictions.

Impact of Number of Concepts: We investigate how the number of concepts in fine-tuned DMs affects auditing performance. As shown in Figure 8, there is no significant difference in performance when more concepts are included, except for a slight degradation when fine-tuning DMs on two celebrity concepts. This suggests that PAIA remains robust across a range of concept complexities in fine-tuned DMs.

Impact of Number of Irrelevant Images: PAIA computes the conditional calibration errors from irrelevant images and utilizes these errors to train the outlier detector. We examine how the number of irrelevant images impacts auditing performance, as shown in Figure 9. While irrelevant images (from concepts that the fine-tuned DMs cannot generate) are generally accessible to the detector, increasing their number can lead to a longer auditing time. By default, we use 200 irrelevant images. The results indicate that using 120 images is sufficient for training an effective detector in PAIA. Adding more irrelevant images beyond this threshold does not significantly improve auditing performance.

6 Evaluation on Real-World DMs.

In this section, we evaluate PAIA on real-world DMs on Civitai, a popular online DM-sharing platform [10].
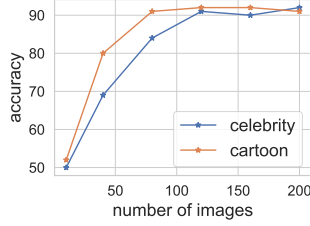
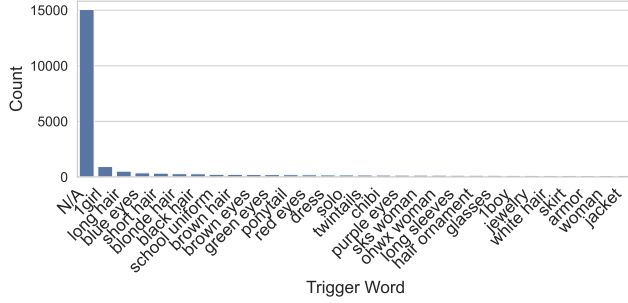**Figure 9: Impact of the number of irrelevant images for training the detector.**

Figure 10: Common trigger words on Civitai. N/A indicates that the user does not provide trigger words.

Table 7: Auditing performance of PAIA on Civitai DMs.

Category	Accuracy	Precision	Recall	F1 Score
Celebrity	93.97%	94.74%	93.10%	93.91%
Cartoon	96.90%	97.22%	96.55%	96.89%
Videogame	95.83%	94.44%	97.40%	95.90%
Movie	92.74%	91.80%	93.85%	92.82%

6.1 Statistic Analysis on Civitai

We first collect and analyze meta from 100, 200 LoRA DMs available on Civitai at the time of collection.

Our analysis reveals that most DMs do not specify trigger words in their metadata (Figure 10). In some cases, "unclaimed" prompts like "ohwx," "sks" are used. These prompts, which are not part of the model vocabulary, are often introduced during fine-tuning to trigger unique concepts. However, the absence of trigger words or the use of uncommon words creates significant challenges for concept auditing, as the intended concepts cannot be easily inferred from the metadata.

To perform a comprehensive evaluation, we conduct an analysis on the models hosted by Civitai. Based on the analysis result (see Section D in Appendix for details), we identify four major topics of concepts: "celebrity," "game," "cartoon," and "movie," and a dominant DM: Stable Diffusion 1.5 (SD1.5). Hence, we download and analyze the LoRA models from these four categories and with Stable Diffusion 1.5 as the base model.

6.2 Model Collection

The model collection process is as follows. First, we search for the LoRA models using the tags provided by each model in Civitai. For the celebrity topic, we use the tag "celebrity," "actress," "real_person," and "actor." For the cartoon topic, we use the tag "cartoon," "manga," and "anime." For the videogame topic, we use the tag of "videogame," "videogames," and "video_game." For the movie topic, we use the tag of "movie." Second, we download the most frequently downloaded LoRA models in the search results. Third, since the model files may stored in different formats (*e.g.*, different parameter names), we

standardize all DMs with consistent formats. Lastly, we manually remove DMs with duplicate concepts (e.g., multiple DMs fine-tuned on the same celebrity or character) and exclude style-related LoRA models that are not concept-specific.

After filtering, we have collected 174 celebrity models, 145 cartoon models, 192 videogame models, and 179 movie models.

6.3 Auditing Performance of PAIA

We evaluate PAIA on the collected LoRA models across the four categories. For evaluation, we use 10 positive samples (images associated with the target concept) and 10 negative samples (images unrelated to the target concept) for each model. Table 7 summarizes the auditing performance of PAIA. PAIA achieves high auditing accuracy, precision, and recall across all categories. Specifically, the cartoon category achieves the best performance with 97% across all metrics, followed closely by the videogame and celebrity categories. The movie category shows slightly lower precision but maintains robust performance overall. These results demonstrate the effectiveness of PAIA in detecting diverse concepts across real-world DMs.

We further analyze the failure cases in the auditing results and observe that a significant portion of these failures arise from concepts that are closely related. Specifically, these include different concepts generated by the same creator or those originating from the same movie or TV show (*e.g.*, Star Wars). Such scenarios introduce shared underlying features in fine-tuned DMs, making it challenging to distinguish between fine-grained variations. We will investigate this challenging scenario and enhance PAIA's fine-grained auditing capabilities in our future work.

7 Related Work

7.1 Inappropriate concept generation in DMs

The creative potential of DMs has raised growing concerns about their ability to generate inappropriate concepts. Current approaches to preventing inappropriate concept generation can be categorized into three main strategies. First, *defensive mechanisms* aim to prevent the generation of certain concepts. For instance, add-on safety filters are employed to detect unsafe input texts or generated images [20]. However, these filters can often be disabled by users after downloading the models, leaving the core DMs capable of generating inappropriate content. Additionally, safety filters often lack generalization and are susceptible to adversarial attacks, undermining their reliability in diverse scenarios [20, 21].

Second, *model refining methods* attempt to modify or remove latent representations of inappropriate concepts through techniques like concept removal and machine unlearning [43–46]. However, these methods are computationally expensive and time-intensive,

making them impractical for real-world applications, particularly for DM-sharing platforms or end users who cannot afford the high costs of retraining models.

Third, *prompt probing techniques* have shown promise in auditing concepts by exploring prompts capable of generating specific concepts [19, 47]. These methods leverage adversarial attacks to optimize prompts, aligning them with target concepts. For example, prompts can be optimized to match the text embeddings of unsafe target prompts filtered by safety mechanisms [19, 48], or to align with image embeddings of a certain concept [18, 49]. However, these techniques rely heavily on accurate prompt probing and robust external detectors to guide the prompt optimization [47]. Unfortunately, external detectors are not typically trained on generated data, particularly for fine-grained concepts, making them vulnerable to concept drift and biased predictions. Our experimental results highlight this limitation: while external detectors perform effectively on real-world images, their reliability degrades significantly when applied to generated content, leading to high false positives.

To the best of our knowledge, there is a critical void of practical and effective tools for large-scale evaluation of DMs.

7.2 Data memorization analysis

Data memorization analysis investigates the extent to which diffusion models memorize their training data, primarily through data extraction attacks and membership inference attacks. For example, Data extraction attacks aim to recover exact data samples from the model's training set [50, 51, 51] while membership inference attacks determine whether a specific data sample was part of the model's training set [52–54]. This line of analysis identifies an exact match of the data samples in the DM's training data, which, if successful, provides strong evidence for concept auditing. However, applying these data memorization approaches in our work is limited, since deriving precise conclusions for individual data samples often requires extensive querying of the DMs and access to a substantial amount of training data for statistical analysis. These demands make such methods impractical and inefficient for large-scale concept auditing.

8 Discussions

8.1 Scope of Concepts

Defining what constitutes a "concept" in DMs remains an open and underexplored problem. Depending on the context, a concept may refer to a broad category, such as "Disney characters," or a highly specific instance, like "Mickey Mouse." This ambiguity exists across related areas such as concept customization [55–57] and concept erasure research [44, 58, 59], where there is no consistent standard for what counts as a distinct or meaningful concept. This ambiguity presents a challenge for concept auditing, as it is difficult to define universal criteria.

In this paper, rather than relying on a rigid definition, we adopt an *example-based approach*: a concept is considered present if the model can generate recognizable outputs that are aligned with a small set of reference examples. This flexible design enables our framework to handle a broad range of concept granularities and types. In our evaluation, we include concepts such as individual

celebrities, full character sets from specific cartoons (*e.g.*, all characters from Pokémon), and object classes from video games. We leave further exploration of formal concept taxonomies and boundaries as future work, particularly toward improving auditing resolution across different concept types.

8.2 Auditing in Multi-LoRA Scenarios

This work mainly focuses on auditing fine-tuned models with a single LoRA module. We investigated the scenarios where multiple concepts are embedded in a single LoRA. Real-world DM deployments also involve multiple LoRAs, which are simultaneously loaded into a DM model. This practice is increasingly common on community platforms, where users compose multiple LoRAs, each capturing different aspects of model behavior such as style, texture, or object identity. The multi-LoRA scenario raises questions about how concept representations may be distributed across different LoRAs, and how interactions between LoRAs might influence a model's generative behavior. We view multi-LoRA auditing as a promising direction for extending the PAIA framework.

8.3 Toward Adversarial Auditing

This work focuses on *non-adversarial auditing*, where fine-tuned DMs are not deliberately optimized to evade auditing. This setting reflects the majority of real-world cases on public platforms, where users unknowingly or negligently publish potentially problematic DMs, often without consistent metadata or clear prompt disclosure. To assess robustness under adversarial conditions, we include a controlled experiment with a simple backdoor-style trigger (Tables 3 and 4), where a concept is activated by an unrelated and abnormal prompt. PAIA remains effective in this setting, highlighting the robustness of its model-centric design.

Nonetheless, backdoor attacks represent an emerging threat to DMs [60–63], where adversaries inject covert patterns into training data or prompts to trigger specific outputs. Existing backdoor detectors primarily focus on defending against image-based triggers [64–66]. Effective and robust prompt-based trigger detection remains an open challenge. Extending PAIA to adversarially evasive fine-tuning introduces a fundamentally different threat model, where the goal is to actively conceal learned concepts. Addressing this challenge would likely require new assumptions, threat models, and detection mechanisms. Moreover, adversarial auditing may conflict with PAIA's scalability goals, as stronger defenses often require expensive or targeted model interrogation. We leave this important but orthogonal direction to future work, and view our current focus as a necessary first step toward practical, large-scale concept auditing.

9 Conclusion

We introduce Prompt-Agnostic Image-Free Detection (PAIA), a model-centric framework for auditing fine-tuned diffusion models without relying on prompt optimization or image-based detection. PAIA combines two key innovations: a *prompt-agnostic design* that analyzes internal model behavior during prompt-insensitive stages of generation, and an *image-free mechanism* based on conditional calibrated error, which compares denoising behavior against the base model to reveal concept learning.

Extensive experiments demonstrate that PAIA consistently outperforms state-of-the-art baselines in both accuracy and efficiency. On 320 benchmark models, PAIA achieves accurate and efficient concept auditing with significantly reduced computation. On 690 real-world models collected from Civitai, it reaches an average accuracy of over 92% across categories including celebrities, cartoons, video games, and movies, highlighting its robustness and real-world applicability.

Our results establish PAIA as the first practical and scalable solution for pre-deployment concept auditing in diffusion models. We hope this work contributes to safer, more transparent model sharing and lays the groundwork for future efforts in responsible generative model management.

References

[1] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):10850–10869, 2023.

[2] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22500–22510, 2023.

[3] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023.

[4] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[5] Diederik P Kingma, Max Welling, et al. Auto-encoding variational bayes, 2013.

[6] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2021.

[7] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021.

[8] automatic1111. stable-diffusion-webui, 2025. https://github.com/AUTOMATIC1111/stable-diffusion-webui, Accessed: January 20, 2025.

[9] Comfy. Comfyui, 2025. https://www.comfy.org/, Accessed: January 20, 2025.

[10] Civitai. Civitai models repository, 2025. https://civitai.com/models, Accessed: January 20, 2025.

[11] Hugging Face. Hugging face models for text-to-image generation, 2025. https://huggingface.co/models?pipeline_tag=text-to-image, Accessed: January 20, 2025.

[12] SeaArt. Seaart model repository, 2025. https://www.seaart.ai/model, Accessed: January 20, 2025.

[13] Yang Zhang, Teoh Tze Tzun, Lim Wei Hern, Haonan Wang, and Kenji Kawaguchi. On copyright risks of text-to-image diffusion models. *arXiv preprint arXiv:2311.12803*, 2023.

[14] 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)*, pages 2187–2204, 2023.

[15] Peifei Zhu, Tsubasa Takahashi, and Hirokatsu Kataoka. Watermark-embedded adversarial examples for copyright protection against diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24420–24430, 2024.

[16] Qingying Hao, Nirav Diwan, Ying Yuan, Giovanni Apruzzese, Mauro Conti, and Gang Wang. It doesn't look like anything to me: using diffusion model to subvert visual phishing detectors. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 3027–3044, 2024.

[17] Yimeng Zhang, Jinghan Jia, Xin Chen, Aochuan Chen, Yihua Zhang, Jiancheng Liu, Ke Ding, and Sijia Liu. To generate or not? safety-driven unlearned diffusion models are still easy to generate unsafe images... for now. *European Conference on Computer Vision (ECCV)*, 2024.

[18] Yuchen Yang, Bo Hui, Haolin Yuan, Neil Gong, and Yinzhi Cao. Sneakyprompt: Jailbreaking text-to-image generative models. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 123–123. IEEE Computer Society, 2024.

[19] Zhi-Yi Chin, Chieh-Ming Jiang, Ching-Chun Huang, Pin-Yu Chen, and Wei-Chen Chiu. Prompting4debugging: Red-teaming text-to-image diffusion models by finding problematic prompts. *ArXiv*, abs/2309.06135, 2023.

[20] Javier Rando, Daniel Paleka, David Lindner, Lennard Heim, and Florian Tramèr. Red-teaming the stable diffusion safety filter. *ArXiv*, abs/2210.04610, 2022.

[21] Yuchen Yang, Bo Hui, Haolin Yuan, Neil Zhenqiang Gong, and Yinzhi Cao. Sneakyprompt: Evaluating robustness of text-to-image generative models' safety

[22] filters. *ArXiv*, abs/2305.12082, 2023.

[22] Gerhard Widmer and Miroslav Kubat. Learning in the presence of concept drift and hidden contexts. *Machine learning*, 23:69–101, 1996.

[23] João Gama, Indrė Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. A survey on concept drift adaptation. *ACM computing surveys (CSUR)*, 46(4):1–37, 2014.

[24] Yuxin Wen, Neel Jain, John Kirchenbauer, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Hard prompts made easy: Gradient-based discrete optimization for prompt tuning and discovery. *Advances in Neural Information Processing Systems*, 36, 2024.

[25] James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, et al. Improving image generation with better captions. *Computer Science. https://cdn. openai. com/papers/dall-e-3. pdf*, 2(3):8, 2023.

[26] Ali Borji. Generated faces in the wild: Quantitative comparison of stable diffusion, midjourney and dall-e 2. *arXiv preprint arXiv:2210.00586*, 2022.

[27] Zihao Fu, Haoran Yang, Anthony Man-Cho So, Wai Lam, Lidong Bing, and Nigel Collier. On the effectiveness of parameter-efficient fine-tuning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 12799–12807, 2023.

[28] Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, et al. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence*, 5(3):220–235, 2023.

[29] Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin A Raffel. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems*, 35:1950–1965, 2022.

[30] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *ArXiv e-prints*, pages arXiv–1607, 2016.

[31] Wentian Zhang, Haozhe Liu, Jinheng Xie, Francesco Faccio, Mike Zheng Shou, and Jürgen Schmidhuber. Cross-attention makes inference cumbersome in text-to-image diffusion models. *arXiv preprint arXiv:2404.02747*, 2024.

[32] Mingyang Yi, Aoxue Li, Yi Xin, and Zhenguo Li. Towards understanding the working mechanism of text-to-image diffusion model. *arXiv preprint arXiv:2405.15330*, 2024.

[33] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 eighth ieee international conference on data mining*, pages 413–422. IEEE, 2008.

[34] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR, 2022.

[35] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*, 2022.

[36] Bairu Hou, Jinghan Jia, Yihua Zhang, Guanhua Zhang, Yang Zhang, Sijia Liu, and Shiyu Chang. Textgrad: Advancing robustness evaluation in nlp by gradient-driven optimization. *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.

[37] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022.

[38] 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*, pages 8748–8763. PMLR, 2021.

[39] Hans-Peter Kriegel, Matthias Schubert, and Arthur Zimek. Angle-based outlier detection in high-dimensional data. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 444–452, 2008.

[40] Jorma Laaksonen and Erkki Oja. Classification with learning k-nearest neighbors. In *Proceedings of international conference on neural networks (ICNN'96)*, volume 3, pages 1480–1483. IEEE, 1996.

[41] Charu C Aggarwal and Charu C Aggarwal. *An introduction to outlier analysis*. Springer, 2017.

[42] Larry M Manevitz and Malik Yousef. One-class svms for document classification. *Journal of machine Learning research*, 2(Dec):139–154, 2001.

[43] Patrick Schramowski, Manuel Brack, Bjorn Deiseroth, and Kristian Kersting. Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22522–22531, 2022.

[44] Rohit Gandikota, Joanna Materzynska, Jaden Fiotto-Kaufman, and David Bau. Erasing concepts from diffusion models. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2426–2436, 2023.

[45] Sanghyun Kim, Seohyeong Jung, Balhae Kim, Moonseok Choi, Jinwoo Shin, and Juho Lee. Towards safe self-distillation of internet-scale text-to-image diffusion models. *ArXiv*, abs/2307.05977, 2023.

[46] Xinfeng Li, Yuchen Yang, Jiangyi Deng, Chen Yan, Yanjiao Chen, Xiaoyu Ji, and Wenyuan Xu. Safegen: Mitigating unsafe content generation in text-to-image models. 2024.

[47] Yimeng Zhang, Jinghan Jia, Xin Chen, Aochuan Chen, Yihua Zhang, Jiancheng Liu, Ke Ding, and Sijia Liu. To generate or not? safety-driven unlearned diffusion models are still easy to generate unsafe images ... for now. *ArXiv*, abs/2310.11868, 2023.

[48] Yijun Yang, Ruiyuan Gao, Xiaosen Wang, Nan Xu, and Qiang Xu. Mma-diffusion: Multimodal attack on diffusion models. *ArXiv*, abs/2311.17516, 2023.

[49] Yuxin Wen, Neel Jain, John Kirchenbauer, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Hard prompts made easy: Gradient-based discrete optimization for prompt tuning and discovery. *ArXiv*, abs/2302.03668, 2023.

[50] Nicolas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwag, Florian Tramer, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion models. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 5253–5270, 2023.

[51] Gowthami Somepalli, Vasu Singla, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Diffusion art or digital forgery? investigating data replication in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6048–6058, 2023.

[52] Jinhao Duan, Fei Kong, Shiqi Wang, Xiaoshuang Shi, and Kaidi Xu. Are diffusion models vulnerable to membership inference attacks? In *International Conference on Machine Learning*, pages 8717–8730. PMLR, 2023.

[53] Tomoya Matsumoto, Takayuki Miura, and Naoto Yanai. Membership inference attacks against diffusion models. In *2023 IEEE Security and Privacy Workshops (SPW)*, pages 77–83. IEEE, 2023.

[54] Zhe Ma, Xuhong Zhang, Qingming Li, Tianyu Du, Wenzhi Chen, Zonghui Wang, and Shouling Ji. Could it be generated? towards practical analysis of memorization in text-to-image diffusion models. *arXiv preprint arXiv:2405.05846*, 2024.

[55] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1931–1941, 2023.

[56] Yan Zeng, Masanori Suganuma, and Takayuki Okatani. An improved method for personalizing diffusion models. *arXiv preprint arXiv:2407.05312*, 2024.

[57] James Seale Smith, Yen-Chang Hsu, Lingyu Zhang, Ting Hua, Zsolt Kira, Yilin Shen, and Hongxia Jin. Continual diffusion: Continual customization of text-to-image diffusion with c-lora. *arXiv preprint arXiv:2304.06027*, 2023.

[58] Shilin Lu, Zilan Wang, Leyang Li, Yanzhu Liu, and Adams Wai-Kin Kong. Mace: Mass concept erasure in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6430–6440, 2024.

[59] Zhili Liu, Kai Chen, Yifan Zhang, Jianhua Han, Lanqing Hong, Hang Xu, Zhenguo Li, Dit-Yan Yeung, and James T Kwok. Implicit concept removal of diffusion models. In *European Conference on Computer Vision*, pages 457–473. Springer, 2024.

[60] Weixin Chen, Dawn Song, and Bo Li. Trojdiff: Trojan attacks on diffusion models with diverse targets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4035–4044, 2023.

[61] 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*, pages 1577–1587, 2023.

[62] Sheng-Yen Chou, Pin-Yu Chen, and Tsung-Yi Ho. How to backdoor diffusion models? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4015–4024, 2023.

[63] Sheng-Yen Chou, Pin-Yu Chen, and Tsung-Yi Ho. Villandiffusion: A unified backdoor attack framework for diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024.

[64] Shengwei An, Sheng-Yen Chou, Kaiyuan Zhang, Qiuling Xu, Guanhong Tao, Guangyu Shen, Siyuan Cheng, Shiqing Ma, Pin-Yu Chen, Tsung-Yi Ho, et al. Elijah: Eliminating backdoors injected in diffusion models via distribution shift. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 10847–10855, 2024.

[65] Yichuan Mo, Hui Huang, Mingjie Li, Ang Li, and Yisen Wang. Terd: A unified framework for safeguarding diffusion models against backdoors. *arXiv preprint arXiv:2409.05294*, 2024.

[66] Yang Sui, Huy Phan, Jinqi Xiao, Tianfang Zhang, Zijie Tang, Cong Shi, Yan Wang, Yingying Chen, and Bo Yuan. Disdet: Exploring detectability of backdoor attack on diffusion models. *arXiv preprint arXiv:2402.02739*, 2024.

A Proof of Lemma 1

We present the full derivation of the gradient of the cross-attention function with respect to the text embedding $\mathbf{p}$.

Lemma 2. *The gradient of the cross-attention function with respect to the text embedding $\mathbf{p}$ is given by:*

$$\frac{\partial Y_i}{\partial \mathbf{p}} = (diag(S_i) - S_i S_i^T)(\frac{1}{\sqrt{d}}(XW_Q)W_K^T)\mathbf{p}W_V + S_i W_V^T, \tag{13}$$

where S denotes the cross-attention map (Eq. 2), S_i the i-th row of S, X the image features calculated by the previous layers, and Y_i the i-th output of the cross-attention layer (Eq. 3).

Proof. The cross-attention function can expressed as

$$Y = \text{softmax}(\frac{QK^T}{\sqrt{d}})V$$

$$= \text{softmax}(\frac{(\mathbf{z}W_Q)(\mathbf{p}W_K)^T}{\sqrt{d}})(\mathbf{p}W_V)$$

Let $A = (XW_Q)(\mathbf{p}W_K)^T/\sqrt{d}$ and $S = \text{softmax}(A)$, where $S_{ij} = \frac{e^{A_{ij}}}{\sum_k e^{A_{ik}}}$. Therefore, we have $Y = S\mathbf{p}W_V$. To compute the gradient of Y with respect to $\mathbf{p}$, we use the chain rule:

$$\frac{\partial Y}{\partial \mathbf{p}} = \frac{\partial S}{\partial \mathbf{p}}(\mathbf{p}W_V) + S\frac{\partial \mathbf{p}W_V}{\partial \mathbf{p}} \tag{14}$$

To calculate the first term, we have the Jacobian matrix of Softmax, which is given by:

$$\frac{\partial S_{ij}}{\partial A_{lk}} = \begin{cases} 0 & i \neq l \\ S_{ij}(\delta_{jk} - S_{jk}) & i = l \end{cases} \tag{15}$$

where δ_{jk} is the Kronecker delta, $\delta_{jk} = 1$ if $j = k$ and 0 otherwise. For simplicity, we denote the Jacobian of S_i with respect to A as:

$$J_i := \frac{\partial S_i}{\partial A} = diag(S_i) - S_i S_i^T. \tag{16}$$

Thus, the gradient in the first term of Eq. 14 becomes:

$$\frac{\partial S_i}{\partial \mathbf{p}} = \frac{\partial S_i}{\partial A}\frac{\partial A}{\partial \mathbf{p}} = J_i(\frac{1}{\sqrt{d}}(XW_Q)W_K^T). \tag{17}$$

By incorporating Eq. 17 into Eq. 14, we have:

$$\frac{\partial Y_i}{\partial \mathbf{p}} = J_i(\frac{1}{\sqrt{d}}(XW_Q)W_K^T)\mathbf{p}W_V + S_i W_V^T$$

□

B Terms used in random prompt generation.

The following set of descriptive keywords was used to augment prompts in the random-prompt evaluation setting. These terms were randomly inserted to simulate realistic prompt noise and diversity.

Table 8: Words used in pseudo prompt generation strategy.

natural lighting	portrait	photorealistic	best quality
realistic	ultra detailed	standing	highres
detailed face	solo	masterpiece	outdoors
film grain	illustration	soft light	raw photo
street	from side	looking at viewer	sitting

C Benchmarking Dataset Samples: Celebrities and Cartoons

We provide example images from the curated datasets used in our benchmark evaluation. Each concept is represented by a small set of reference images used during fine-tuning and for auditing evaluation.

Figure 11: Example concept images from the Celebrity dataset used in benchmarking.

Figure 12: Example concept images from the Cartoon character dataset used in benchmarking.

D Civitai model analysis.

To support real-world evaluation, we analyze the diffusion models hosted on Civitai, one of the largest community platforms for sharing LoRA-based fine-tuned models.

Figure 13 shows the distribution of commonly used tags on Civitai. Among these, the most frequent concept categories include "celebrity," "game," and "cartoon." Based on this trend, our evaluation focuses on LoRA models falling within these categories. Additionally, we include a "movie" category to capture models, which overlaps with tags such as "character" and "actress".

Figure 14 shows that Stable Diffusion 1.5 (SD1.5) is the most commonly used base model for LoRA fine-tuning. Therefore, our real-world evaluation focuses primarily on LoRA models adapted from SD1.5.
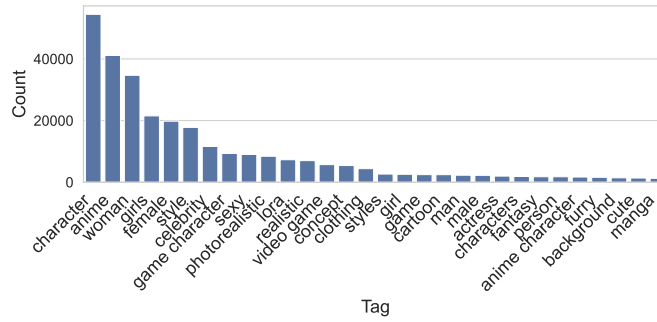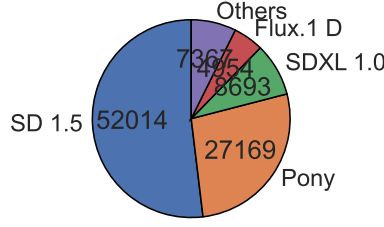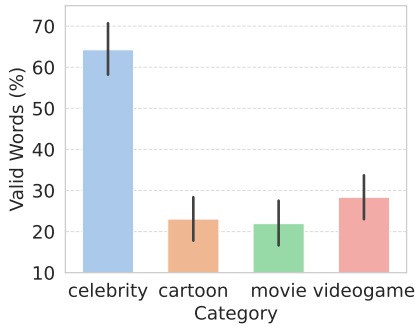

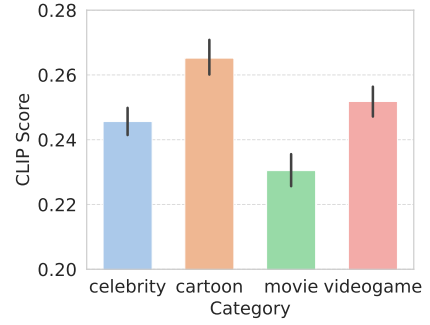Figure 13: Common tags in models uploaded to Civitai.

Figure 14: Distribution of base models used for LoRA fine-tuning on Civitai.

E   Trigger Quality Analysis on Civitai Models

To better understand the challenges in prompt-based auditing, we analyze the trigger words associated with LoRA models uploaded to Civitai. While many models provide trigger prompts, the quality and alignment of these prompts vary significantly.

(a) Percentage of valid English words used in trigger words.

(b) CLIP similarity between generated images and their associated trigger words.

Figure 15: Analysis of trigger words that are associated with uploaded DMs over four categories on Civitai.

Figure 15a shows that a large portion of trigger words used by creators are not valid English words. These nonstandard or synthetically generated tokens are common in community models, making them difficult to probe or interpret.

Figure 15b further highlights the issue by quantifying the semantic alignment between trigger words and their corresponding images. We use CLIP similarity scores to measure how well a trigger word matches the content of its sample image. The observed scores reveal that many triggers exhibit weak semantic correspondence with the generated outputs, indicating that even when a trigger is known, it may not meaningfully reflect the learned concept.

Together, these findings underscore a key obstacle in prompt-based auditing—**trigger uncertainty**, which arises from the large, discrete nature of the prompt space and the inherent ambiguity of natural language. Optimization-based probing methods (*e.g.*, adversarial or reinforcement learning) often fail to reliably discover effective prompts, especially in the presence of low-quality or misleading triggers. These issues further motivate our model-centric, prompt-agnostic auditing framework.