

Stabilizing Self-Consuming Diffusion Models with Latent Space Filtering

Zhongteng Cai¹, Yaxuan Wang², Yang Liu², Xueru Zhang¹

¹ The Ohio State University

² University of California, Santa Cruz

cai.1125@osu.edu, ywan1225@ucsc.edu, yangliu@ucsc.edu, zhang.12807@osu.edu

Abstract

As synthetic data proliferates across the Internet, it is often reused to train successive generations of generative models. This creates a “self-consuming loop” that can lead to training instability or *model collapse*. Common strategies to address the issue—such as accumulating historical training data or injecting fresh real data—either increase computational cost or require expensive human annotation. In this paper, we empirically analyze the latent space dynamics of self-consuming diffusion models and observe that the low-dimensional structure of latent representations extracted from synthetic data degrade over generations. Based on this insight, we propose *Latent Space Filtering* (LSF), a novel approach that mitigates model collapse by filtering out less realistic synthetic data from mixed datasets. Theoretically, we present a framework that connects latent space degradation to empirical observations. Experimentally, we show that LSF consistently outperforms existing baselines across multiple real-world datasets, effectively mitigating model collapse without increasing training cost or relying on human annotation.

1 Introduction

Diffusion models have emerged as a leading class of generative models in computer vision, achieving state-of-the-art performance in image generation tasks (Dhariwal and Nichol 2021). By learning to reverse a noise-injection process, these models can synthesize high-quality samples (Ho, Jain, and Abbeel 2020; Song and Ermon 2019). Compared to earlier approaches such as variational autoencoders (VAEs) (Chen et al. 2024) and generative adversarial networks (GANs) (Goodfellow et al. 2020), diffusion models offer greater training stability, improved sample fidelity (Dhariwal and Nichol 2021), and broader applicability across diverse domains.

Modern diffusion models are typically trained on large-scale datasets scraped from the Internet. However, as synthetic data proliferates online, it inevitably becomes part of the training data for future generations of models, creating a “self-consuming training loop.” Recent studies have shown that such self-consuming loops may lead to model collapse (Shumailov et al. 2023; Gerstgrasser et al. 2024; Alemohammad et al. 2024a; Shumailov et al. 2024; Wei and Zhang 2025), training instability (Bertrand et al. 2024), and potential

bias amplification (Taori and Hashimoto 2023; Wyllie, Shumailov, and Papernot 2024; Xie and Zhang 2024). To mitigate these issues, several strategies have been proposed, including: (i) accumulating all historical (real and synthetic) samples into the training dataset (Gerstgrasser et al. 2024; Wang et al. 2024); (ii) injecting sufficient fresh real data during each generation cycle (Alemohammad et al. 2024a; Kazdan et al. 2024; Fu et al. 2025); (iii) modifying the training process in each round—for example, by utilizing prior knowledge of the real distribution (Gillman et al. 2024), merging sequentially trained models (Alemohammad et al. 2024b; Yuan et al. 2025), or leveraging feedback on synthetic data (Feng et al. 2025). However, these approaches often incur substantial storage and computational costs, require expensive and reliable human annotations, or are limited to single-step retraining rather than the full multi-step self-consuming process.

In this work, we address the above limitations and propose an alternative strategy to mitigate model collapse without introducing additional storage overhead, training cost, or reliance on human annotations. Unlike existing works that study self-consuming generative models in input space, we shift the focus to latent space, aiming to understand the causes of model collapse through the lens of latent representations. The idea is to first investigate how the representations of model-generated data evolve in latent space across successive self-consuming training loops. Based on this insight, we develop a filtering mechanism that mitigates model collapse by identifying and removing less realistic data from the mixed training set at each round of the self-consuming process.

Specifically, we use diffusion models to extract latent representations for real and synthetic images sampled at each round of the self-consuming training loop (Li et al. 2025). For each generation, we compute Orthogonal Low-rank Embedding (OLE) scores (Lezama et al. 2018) of the latent representations, which quantify the orthogonality of class-specific subspaces and reflect the underlying low-dimensional structure. Our key observation is that the low-dimensional structure of latent representations from synthetic images gradually degenerates as self-consuming training progresses, whereas real images tend to exhibit more orthogonal subspaces. Motivated by this observation, we introduce *Latent Space Filtering* (LSF), a method that identifies and filters synthetic images with degraded latent representations before they are used in further training. Specifically, we train a representation prob-

ing classifier for real images based on their latent embeddings. The confidence score from this classifier serves as a proxy for how well a data sample aligns with the low-dimensional manifold of real data. Samples with low alignment are discarded, preventing them from introducing additional noise into the model’s learned latent subspaces. In addition to the algorithmic development, we establish a theoretical foundation for the proposed method, which explains how the degeneration of low-dimensional structures correlates with changes in both the OLE scores and the probing classifier’s confidence scores. Our main contributions are summarized below.¹

- In Section 2, we formalize the problem of studying self-consuming diffusion models in latent space. To the best of our knowledge, this is the first work to investigate how the low-dimensional structure of latent representations evolves across the self-consuming training loop.
- In Section 3, we analyze how the low-dimensional structure of latent representations—measured by Orthogonal Low-rank Embedding (OLE) scores—evolves across successive generations of self-consuming diffusion models.
- Motivated by empirical observations, Section 4 proposes *Latent Space Filtering* (LSF), a practical approach that mitigates model collapse by filtering low-quality synthetic data based on subspace alignment, without requiring additional real data collection or increased training cost.
- We establish a theoretical framework to justify both empirical observations and the proposed method. This framework explains how the degeneration of low-dimensional structures correlates with changes in OLE scores (Theorem 1) and in the confidence scores of the probing classifier used to assess subspace alignment (Theorem 2).
- In Section 5, we extensively evaluate the proposed method across multiple real-world datasets². The results demonstrate that LSF effectively mitigates model collapse and outperforms existing baselines, without the need for a growing training set or additional curated data.

2 Problem Formulation

Consider a platform that iteratively trains diffusion models using a sequence of training datasets. Let $f^{(k)}$ denote the diffusion model trained on dataset $\mathcal{D}^{(k)}$ during the k -th round of the iterative retraining loop, and let $\widehat{\mathcal{D}}^{(k+1)}$ be the synthetic dataset generated by $f^{(k)}$. Without loss of generality, we assume the initial dataset $\mathcal{D}^{(0)}$ used to train the first model $f^{(0)}$ consists of real data. We consider two variants of the self-consuming training loop (Alemohammad et al. 2024a):

1. **Pure synthetic loop:** Only the synthetic data generated by the most recent model is used to train the next model, i.e., $\mathcal{D}^{(k)} = \widehat{\mathcal{D}}^{(k)}, \forall k$.
2. **Accumulation loop:** All previously used training data, including the initial real dataset $\mathcal{D}^{(0)}$ and synthetic samples from earlier generations, are accumulated to train the

¹More related work is discussed in Appendix D. Conclusions and limitations are provided in Appendix E.

²Source code available at <https://github.com/osu-srml/Latent-Space-Filtering>

next model, i.e., $\mathcal{D}^{(k)} := \mathcal{D}^{(k-1)} \cup \widehat{\mathcal{D}}^{(k)} = \mathcal{D}^{(0)} \cup \widehat{\mathcal{D}}^{(1)} \cup \dots \cup \widehat{\mathcal{D}}^{(k)}, \forall k$.

In practice and the experiments presented in Section 5, we may apply additional processing steps to $\mathcal{D}^{(k)}$, such as re-sampling or filtering, before training $f^{(k)}$.

Background: Diffusion Model. We use Denoising Diffusion Probabilistic Models (DDPM) (Ho, Jain, and Abbeel 2020) as an example to illustrate the training of $f^{(k)}$ in each round of the retraining loop, though our analysis and method are broadly applicable beyond this specific architecture.

Training a DDPM consists of two stages: a *forward process* that gradually adds noise to the data, and a *reverse process* that learns to remove the noise. In the forward process, Gaussian noise is iteratively added to a clean image $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ over T steps, producing a sequence of increasingly noisy images $\{\mathbf{x}_1, \dots, \mathbf{x}_T\}$. The conditional distribution at each timestep is defined as:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}),$$

where β_t is a variance schedule. As $T \rightarrow \infty$, the image \mathbf{x}_T approaches pure Gaussian noise. The reverse process learns to denoise $\epsilon_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ back to \mathbf{x}_0 using a model $\widehat{\epsilon}_\theta$ trained to predict the noise at each step. Specifically, the model is optimized by minimizing the simplified Evidence Lower Bound (ELBO) objective: $\nabla_\theta \|\widehat{\epsilon}_\theta(\mathbf{x}_t, t) - \epsilon_0\|^2$. At generation time, denoising proceeds from $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to \mathbf{x}_0 using the learned $\widehat{\epsilon}_\theta$ according to the following:

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \widehat{\epsilon}_\theta(\mathbf{x}_t, t) \right) + \sigma_q(t) \mathbf{z},$$

where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$, $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and $\sigma_q(t) = \sqrt{\frac{(1 - \alpha_t)(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}}$. Song and Ermon (2019) propose a related framework based on score matching, where a neural network is trained to approximate the score function $\mathbf{s}_\theta(\mathbf{x}) \approx \nabla_{\mathbf{x}} \log q(\mathbf{x})$. Using Langevin dynamics (Welling and Teh 2011), one can then generate samples that follow the target distribution via the learned score function. This framework connects to DDPM through the following approximation: $\mathbf{s}_\theta(\mathbf{x}_t, t) \approx -\frac{\widehat{\epsilon}_\theta(\mathbf{x}_t, t)}{\sqrt{1 - \alpha_t}}$.

In practice, diffusion models typically employ U-Net architectures (Ronneberger, Fischer, and Brox 2015), which encode features through a series of downsampling layers and reconstructs the image via upsampling. We extract latent representations from U-Net’s bottleneck layer, which has proven effective for various downstream tasks such as controllable image editing (Kwon, Jeong, and Uh 2023; Li et al. 2025).

Low-dimensional structure of latent representations under self-consuming loop. Existing research on self-consuming generative models has primarily focused on the input space, while the impact on latent representations remains largely unexplored. Understanding how latent representations and their underlying low-dimensional structure evolve during self-consuming training may offer valuable insights for developing strategies to mitigate model collapse.

Specifically, for the k -th round of iterative retraining loop, let the denoising model of the diffusion model $f^{(k)}$ be denoted as $\widehat{\epsilon}_\theta^{(k)}(\cdot) = d^{(k)} \circ e^{(k)}(\cdot)$, which is formed by an

encoder $e^{(k)}$ and a decoder $d^{(k)}$. Since $e^{(0)}$ is the encoder of the initial model $f^{(0)}$ trained solely on *real* data $\mathcal{D}^{(0)}$, we consider the output of $e^{(0)}$ as the **latent representation** of a data sample in this work. For a sample $\mathbf{x}^{(k)}$ generated by $f^{(k)}$ at k -th round of self-consuming training loop, we denote the latent representation at denoising timestep t as $\mathbf{h}_t^{(k)} = e^{(0)}(\mathbf{x}^{(k)}, t)$, following prior works on representation extraction via diffusion models (Li et al. 2025).

Objectives. In this work, we focus on self-consuming diffusion models and aim to leverage the low-dimensional structure of latent representations to better understand the causes of model collapse and develop effective mitigation strategies. Our goals are to: 1) examine how the low-dimensional structure of latent representations $\mathbf{h}_t^{(k)}$ evolves throughout the self-consuming training loop; and 2) based on these insights, develop a novel mechanism to mitigate model collapse by filtering out less realistic samples, which compared to prior methods, without incurring additional training costs or relying on human annotations.

3 Evolution of Low-Dimensional Structure of Latent Representations

Empirical observations. For a batch of N images $\mathcal{B}^{(k)} = \{(\mathbf{x}_1^{(k)}, y_1^{(k)}), \dots, (\mathbf{x}_N^{(k)}, y_N^{(k)})\}$ sampled from dataset $\mathcal{D}^{(k)}$ at k -th round of self-consuming loop,³ the latent representations extracted using $e^{(0)}$ in matrix form are given by:

$$\mathbf{M}_t^{(k)} = \begin{bmatrix} \mathbf{h}_{1,t}^{(k)} & \mathbf{h}_{2,t}^{(k)} & \dots & \mathbf{h}_{N,t}^{(k)} \end{bmatrix},$$

where $\mathbf{h}_{i,t}^{(k)} = e^{(0)}(\mathbf{x}_i^{(k)}, t)$ (1)

Similarly, we define the *class-specific* latent representations $\mathbf{M}_{c,t}^{(k)}$ for all samples $\mathcal{B}_c^{(k)} = \{(\mathbf{x}, y) \in \mathcal{B}^{(k)} \mid y = c\}$ belonging to class $c \in \mathcal{C}$. Given (class-specific) latent representations $\mathbf{M}_t^{(k)}$ and $\{\mathbf{M}_{c,t}^{(k)}\}_{c \in \mathcal{C}}$, we measure the low-dimensional structure and class separability using the Orthogonal Low-rank Embedding (OLE) criterion (Lezama et al. 2018), which captures the degree of orthogonality between class-specific subspaces and is defined as:

$$\text{OLE}_t^{(k)} := \sum_{c \in \mathcal{C}} \|\mathbf{M}_{c,t}^{(k)}\|_* - \|\mathbf{M}_t^{(k)}\|_*,$$

where $\|\cdot\|_*$ is the nuclear norm (sum of the singular values). A lower OLE value indicates that the latent representations exhibit stronger intra-class low-rank structure and greater inter-class orthogonality, thereby reflecting a more structured and separable latent space. The OLE criterion has been widely used in prior work to enhance feature discriminability in classification tasks (Lezama et al. 2018; Huang, Qiu, and Sapiro 2025). In our setting, increasing OLE values suggest a trend toward a degenerate latent space, where class-specific subspaces lose structure and become less distinguishable.

³The class label can either be the ground-truth annotation or a group label assigned by an unsupervised clustering algorithm. Unsupervised clustering enables us to explore low-dimensional structures beyond predefined classes, as diffusion models can capture factorized features that may not align with standard semantic categories.




Figure 1: OLE values of latent representations extracted by a fixed diffusion model across generations and denoising timesteps. Each curve corresponds to a generation within a pure synthetic self-consuming loop. When conditioned on timestep, OLE values increase with generation, indicating progressive structural degradation of the latent space. When conditioned on generation, OLE exhibits a U-shaped trend over denoising timesteps.

We first consider a pure synthetic self-consuming loop and train diffusion models for 6 generations, with models fine-tuned on 1,000 images for 3 epochs in each generation. For each image, we assign a label using sparse subspace clustering by orthogonal matching pursuit (SSC-OMP) algorithm (You, Robinson, and Vidal 2016). We compute OLE values for 10 batches of real and synthetic data. The results for MNIST (LeCun, Cortes, and Burges 2010) and CelebA (Liu et al. 2015) are shown in Figure 1. Conditioned on generation, we observe that OLE initially decreases and then increases with the denoising timestep; this trend is consistent with recent findings that representation quality exhibits a unimodal trajectory (Li et al. 2025). Lower OLE values are associated with greater feature separability, higher representation quality, and improved performance on downstream tasks. More importantly, when conditioned on the same timestep, OLE values increase as the self-consuming generation progresses. This suggests that the latent representations become less structured and increasingly entangled over time.

The above observations are expected. Intuitively, latent features extracted by the diffusion model act as intermediaries that guide the generation of high-quality images. However, in self-consuming loops, synthetic data gradually deviates from the real data distribution. As a result, the learned low-dimensional structure diverges from the true data manifold, becoming less informative over time and ultimately leading to a decline in the quality of generated samples.

Theoretical analysis. The formulation of OLE is motivated by geometric intuition, aiming to quantify how well class-specific feature subspaces are aligned or separated. Prior work has shown that OLE is always non-negative and equals zero when the subspaces are perfectly orthogonal (Lezama et al. 2018). However, the quantitative relationship between OLE and the degree of subspace orthogonality remains unexplored. In what follows, we address this gap by developing a theoretical framework that characterizes the connection between latent space degeneration and the dynamics of OLE. For clarity of presentation, we omit the diffusion timestep t

and generation index k in the notations.

Given latent representation \mathbf{h}_i of data sample (\mathbf{x}_i, y_i) generated by encoder $e^{(0)}$ (Eqn. (1)), we assume data are from two classes $\mathcal{C} = \{0, 1\}$ and latent representations $\mathbf{h}_i \in \mathbb{R}^{d \times 1}$ for each class $c \in \{0, 1\}$ are drawn from noisy low-rank Gaussian distributions. Mathematically, we model conditional distribution of $\mathbf{h}_i \mid y_i = c$ as zero-mean Gaussian $\mathcal{N}(\mathbf{0}, \Sigma_c)$ with covariance $\Sigma_c = \mathbf{U}_c \mathbf{U}_c^\top + \sigma^2 \mathbf{I}_d$. Here, $\mathbf{U}_c \in \mathbb{R}^{d \times r}$ is an orthonormal basis matrix representing the low-dimensional subspace of class c with rank $r < d$, and $\sigma^2 \mathbf{I}_d$ represents isotropic noise.

Let matrix $\mathbf{M}_c \in \mathbb{R}^{d \times n}$ be the *class-specific* latent representations for n samples from class c , with each column the representation of one sample. Let $\mathbf{M} = [\mathbf{M}_0 \ \mathbf{M}_1] \in \mathbb{R}^{d \times 2n}$ denote the concatenated matrix of a batch of samples from two classes. The OLE for this sample matrix becomes:

$$\text{OLE}(\mathbf{M}_0, \mathbf{M}_1) = \|\mathbf{M}_0\|_* + \|\mathbf{M}_1\|_* - \|\mathbf{M}\|_*,$$

where $\|\cdot\|_*$ denotes the nuclear norm. For any column unit vectors \mathbf{u}_i of \mathbf{U}_0 and \mathbf{v}_j in \mathbf{U}_1 , denote their angle as θ_{ij} , i.e., $\mathbf{u}_i^\top \mathbf{v}_j = \cos \theta_{ij}$, $i, j \in [r]$. Suppose that the cosine similarity of each pair of unit basis can be bounded by:

$$\cos \tilde{\theta} \leq \cos \theta_{ij} \leq \ell \cos \tilde{\theta},$$

where ℓ measures the range of pairwise cosine similarities, $\tilde{\theta} \in [0, \frac{\pi}{2}]$ is the largest pairwise angle. Theorem 1 below provides a lower bound on the expected value of OLE, characterized by the degree of orthogonality between subspaces.

Theorem 1 (Lower bound of OLE). *For all $\theta \in [0, \frac{\pi}{2}]$, when $r > 2n$, the expected OLE satisfies the following:*

$$\mathbb{E}[\text{OLE}(\mathbf{M}_0, \mathbf{M}_1)] \geq C_1 - C_2 \cdot \phi(\tilde{\theta}),$$

where

$$\phi(\tilde{\theta}) := \sqrt{2n} \cdot \sqrt{\ell \cos \tilde{\theta}} + \sqrt{2n(2n-1)} \cdot \sqrt{1 - \cos \tilde{\theta}},$$

$$C_1 := 2\gamma(rn) - 2\sqrt{n}\gamma(dn) - \sqrt{2n} \cdot \gamma(2dn),$$

$$C_2 := 2\gamma(rn),$$

and $\gamma(n) := \sqrt{2} \cdot \frac{\Gamma(\frac{n+1}{2})}{\Gamma(\frac{n}{2})}$, with $\Gamma(\cdot)$ being Gamma function.

The proof of Theorem 1 is provided in Appendix A. This result offers a theoretical explanation for the empirical trend observed in Figure 1, where OLE values increase across successive generations. More specifically, assuming that as the self-consuming process progresses, the quality of the generated data decreases and the learned latent subspaces become less orthogonal (i.e., θ between subspaces $\text{span}(\mathbf{U}_0)$ and $\text{span}(\mathbf{U}_1)$ decreases), Theorem 1 shows that as θ decreases, the lower bound on the expected OLE increases.

Remark 1. *Since $\phi(\theta)$ in Theorem 1 is decreasing as θ decreases over $[0, \arccos \frac{1}{2n}]$, the lower bound of OLE increases as the subspaces $\mathbf{U}_0, \mathbf{U}_1$ become less orthogonal. When $n \rightarrow \infty$, we have $\arccos \frac{1}{2n} \rightarrow \frac{\pi}{2}$, indicating that the conclusion in Theorem 1 holds for most $\theta \in [0, \frac{\pi}{2}]$.*

Mitigate model collapse through filtering. Previous studies have suggested that removing low-quality synthetic samples




Figure 2: Correlation between the average generation number of each batch and its corresponding OLE score, computed on the accumulated CelebA dataset. Batches with lower average generation numbers (i.e., containing more realistic images) tend to have lower OLE values. However, the correlation is not strong enough to enable reliable filtering based solely on batch-level OLE.




Figure 3: Correlation between the OLE and average confidence scores of each batch, computed on the real and synthetic MNIST datasets produced at different generations. Batches with lower OLE (i.e., greater feature separability) tend to have higher confidence scores, indicating that we can use confidence score as an individual-level proxy of batch-level OLE score.

from the training data can be more effective in stabilizing retraining than adding high-quality data (Kazdan et al. 2024). However, identifying synthetic images directly within a mixture of real and synthetic data is challenging (Corvi et al. 2023; Feng et al. 2025). The observations in Figure 1 suggest a potential direction:

Can synthetic images be identified and filtered based on the quality of their latent representations?

A straightforward idea is to compute the OLE score for each batch $\mathcal{B}^{(k)}$ and discard those whose scores exceed a certain threshold. While this approach is intuitive, it becomes less effective in realistic scenarios where real and synthetic images are mixed within each batch. Unlike in our controlled experiments where each batch contained either only real or only synthetic data (as shown in Figure 1), mixed batches dilute the signal that OLE captures.

To assess the feasibility of OLE-based filtering in such mixed scenarios, we combine the accumulated datasets $\mathcal{D}^{(0)}, \dots, \mathcal{D}^{(5)}$ from the Accumulation Loop and compute the OLE score for each batch. Each image is tagged with the generation index from which it was sampled, allowing us to calculate the average generation number per batch. A lower average generation number indicates that the batch contains more samples from earlier stages of self-consuming training, which are less affected by model collapse and therefore more likely to be of higher quality. As shown in Figure 2, there is a positive correlation between batch OLE scores and the average generation number. Since earlier generations contain more realistic samples, this correlation supports the hypothesis that lower OLE scores reflect higher-quality latent representations. However, the correlation is relatively weak,

and filtering data batches based on OLE scores is not a reliable strategy—particularly when each batch contains a mix of real and synthetic images. Indeed, as shown in Figure 4, OLE-based filtering does not yield significant improvement over random sampling in stabilizing self-consuming training. We thus introduce an alternative proxy of OLE scores with a better filtering algorithm, as detailed in Section 4.

4 Proposed Method: Latent Space Filtering

Section 3 shows that OLE scores computed over data batches cannot serve as a reliable criterion for filtering synthetic data, particularly when each batch contains a mix of real and synthetic samples. We therefore introduce an alternative representation-based filtering mechanism, which is related to OLE but more effective, as detailed below.

Motivation. As shown in Figure 1, OLE exhibits a U-shaped trend over diffusion timestep t when conditioned on generation number k , and increases with generation k when conditioned on timestep t . Interestingly, a similar U-shaped pattern was observed in (Li et al. 2025), which demonstrates that the *representation quality* of diffusion models during the denoising process follows a unimodal trajectory. We thus hypothesize that representation quality—measured by the accuracy of a probing classifier (a simple, often linear classifier applied to latent representations to assess their informativeness)—also reflects the underlying low-dimensional structure, similar to what is captured by OLE scores.

Following this idea, we propose a *representation-based* filtering approach. Assuming that a probing classifier trained on the latent representations of *real* images captures the underlying low-dimensional structure, we can then use the classifier’s confidence scores to measure the degree of misalignment between a sample’s representation and those of real data. Specifically, we first extract latent representations of real images using a well-trained diffusion model, then train a softmax regression model based on these representations. Let the output logit for class c be defined as $o_c(\mathbf{x}) := \mathbf{w}_c^\top \mathbf{x} + \mathbf{b}_c$. Then, for an image \mathbf{x} with label y , we compute its **confidence score** regarding the correct class by softmax function:

$$\xi(\mathbf{x}, y) := \frac{\exp\{o_y(\mathbf{x})\}}{\sum_{c \in \mathcal{C}} \exp\{o_c(\mathbf{x})\}}. \quad (2)$$

We hypothesize that images with higher confidence scores are more likely to resemble real data and can therefore serve as a reliable filtering criterion. To verify this hypothesis, we conduct experiments on the MNIST dataset, with results shown in Figure 4.

Figure 4a presents box-and-whisker plots of confidence scores across different generations⁴. The results reveal a clear correlation: **images from later generations tend to have lower confidence scores**. This suggests that confidence scores can serve as a more fine-grained filtering criterion compared to batch-level OLE scores. Furthermore, Figure 3

⁴Each box spans from the first quartile (Q1) to the third quartile (Q3) of the data, with whiskers extending to 1.5 times the interquartile range (IQR). Data points beyond the whiskers are shown as fliers, indicating outliers.

Algorithm 1: Latent Space Filtering


Input: Pre-trained diffusion model $f^{(0)}$, training budget N , self-consuming loop K .

- 1: Train a probing classifier with $f^{(0)}$ and real images $\mathcal{D}^{(0)}$
 - 2: **for** $k = 1$ to K **do**
 - 3: Sample a new dataset $\widehat{\mathcal{D}}^{(k)}$ from model $f^{(k-1)}$.
 - 4: Construct a dataset $\mathcal{D}^{(k)}$ according to the type of the self-consuming loop.
 - 5: Compute the confidence score $\xi(\mathbf{x}_i, y_i)$ for each image $(\mathbf{x}_i, y_i) \in \mathcal{D}^{(k)}$ by Eqn. (2).
 - 6: Select top- N images with highest scores to construct a new set $\underline{\mathcal{D}}^{(k)}$ with training budget N .
 - 7: Train a new diffusion model $f^{(k)}$ using dataset $\underline{\mathcal{D}}^{(k)}$.
 - 8: **end for**
 - 9: **Output:** $f^{(K)}$
-


shows a **positive correlation between low OLE scores and high average confidence scores across batches**, suggesting that confidence scores can serve as a viable proxy for OLE scores. Unlike OLE, which operates at the batch level, confidence scores are computed per image, allowing us to circumvent the previously mentioned issue of mixed real and synthetic samples within the same batch.

To evaluate the effectiveness of confidence-based filtering, we compare it with random sampling in Figures 4b and 4c. For each generation number k , we construct an accumulated dataset by concatenating the real dataset with a series of synthetic datasets $\widehat{\mathcal{D}}^{(1)}, \dots, \widehat{\mathcal{D}}^{(k)}$, where the size of these datasets increases from 1,000 to 5,000 images, as indicated in the figures. Using a pre-trained probing classifier, we score all images and select the top 1,000 with the highest confidence. Figure 4b reports the average generation number of the filtered images: lower values indicate samples that are closer to the real distribution. Since real images play a more critical role in mitigating model collapse in self-consuming loops (Alemohammad et al. 2024a; Fu et al. 2025, 2024), Figure 4c further reports the proportion of real images in the filtered set. The results demonstrate that **filtering using the probing classifier yields both a lower average generation number and a higher fraction of real images** compared to random sampling. Furthermore, as the size of the accumulated dataset increases, the filtering becomes more effective because the larger pool of real images results in the filtered fixed-size dataset containing more real images.


In Algorithm 1, we summarize the complete procedure for applying the proposed filtering method to self-consuming accumulation training loop. It is worth noting that access to real images is only required at the start of the self-consuming loops. Afterward, there is no need to distinguish real images from the synthetic dataset in the accumulated dataset. Additionally, even when images are unlabeled, labeling or captioning each image can be easier than detecting synthetic images or acquiring new data (Sharifzadeh et al. 2024; Lu et al. 2023). By retaining as many real and realistic synthetic images as possible in the filtered dataset, we aim to prevent model collapse while adhering to a fixed training budget.



(a) Generations vs. Confidence scores



(b) Average generation after filtering



(c) Ratio of real images after filtering

Figure 4: The effect of using confidence scores for filtering on MNIST dataset. (a) shows the generation number vs. distribution of confidence scores of images sampled at this generation. Higher generation is correlated with lower confidence score, indicating that they are less aligned with real images, hence providing a strong signal for filtering out unrealistic synthetic images. (b) shows the average generation number within the dataset filtered by the confidence score. (c) Our method can also preserve more real images in the training set. Using larger accumulated dataset for filtering can increase the ratio of selected real images. Compared with randomly sampling from the accumulated dataset, our filtering method can preserve a distribution closer to real images.

Theoretical analysis. Following the same setting as in Theorem 1, we next theoretically characterize the relation between confidence scores and lower-dimensional structure of latent representations. For samples from two classes $\mathcal{C} = \{0, 1\}$, let the log-likelihood ratio for a sample representation \mathbf{h} be defined as $g(\mathbf{h}) = \log \frac{\Pr[\mathbf{h}|y=1]}{\Pr[\mathbf{h}|y=0]}$. Then, the posterior probability $\Pr[y = 1 | \mathbf{h}]$ assigned by the Bayes optimal classifier is given by $\varsigma(g(\mathbf{h}))$, where $\varsigma(\cdot)$ denotes the Sigmoid function. We define the expected confidence score for class 1, under the distribution of class-1 samples, as:

$$\xi(\theta) := \mathbb{E}_{\mathbf{h}|y=1 \sim \mathcal{N}(\mathbf{0}, \Sigma_1)}[\varsigma(g(\mathbf{h}))],$$

where $\theta \in [0, \frac{\pi}{2}]$ is the angle between subspaces spanned by \mathbf{U}_0 and \mathbf{U}_1 . Similar to Theorem 1, the following theorem characterizes the relationship between the confidence score and the degree of orthogonality between subspaces.

Theorem 2 (Upper bound of confidence score). *The expected confidence score satisfies the following:*

$$\xi(\theta) \leq \frac{1}{2\sigma^2(\sigma^2 + 1)} \varsigma(r \sin^2 \theta),$$

where r is the rank of \mathbf{U}_0 and \mathbf{U}_1 . The proof is provided in Appendix B. Note that the upper bound is monotonically increasing in θ . As $\theta \rightarrow 0$, the subspaces $\text{span}(\mathbf{U}_0)$ and $\text{span}(\mathbf{U}_1)$ become more aligned, making the two classes less distinguishable. Consequently, the upper bound of the expected confidence score decreases, reflecting the reduced separability between classes.

5 Experiments

Datasets. We evaluate our approach by training DDPM diffusion models (Ho, Jain, and Abbeel 2020) for conditional image synthesis (Dhariwal and Nichol 2021) on three real-world datasets: (i) **MNIST** (LeCun, Cortes, and Burges 2010): A dataset of 60,000 grayscale handwritten digits across 10 classes, with image size 28×28 . (ii) **CIFAR-10** (Krizhevsky,

Hinton et al. 2009): A collection of 60,000 32×32 RGB images spanning 10 object categories. (iii) **CelebA** (Liu et al. 2015): A dataset of 200,000 celebrity face images, each annotated with 40 binary attributes⁵.

Baselines. For each dataset, we first train a baseline model on the full dataset. We then iteratively fine-tune the model for five self-consuming generations. Our method, *ACU-LSF*, is compared with the following five baselines:

1. **SYN**: Each generation is trained purely on synthetic data generated by the previous model.
2. **SYN-ADD** (Alemohammad et al. 2024a): A hybrid setting where each generation’s training data consists of 70% synthetic images and 30% newly added real images. The 30% proportion is chosen to approximate the average percentage of real images present in the accumulated datasets over five generations.
3. **ACU** (Gerstgrasser et al. 2024): Historical real and synthetic data are accumulated across generations, resulting in a linearly growing training set. A subset of real data is sampled for accumulation in each generation, matching the size of newly generated synthetic data.
4. **ACUR** (Kazdan et al. 2024): A fixed-size subset is randomly sampled from the accumulated dataset used in ACU. This setting evaluates the impact of maintaining a fixed training budget.
5. **ACUR-SIMS** (Alemohammad et al. 2024b): This method extrapolates the score function to enhance fidelity. At generation k , the score function used for sampling is:

$$\mathbf{s}_\theta^{(k)}(\mathbf{x}_t, t) = (1 + \omega) \mathbf{s}_\theta^{(k-1)}(\mathbf{x}_t, t) - \omega \widehat{\mathbf{s}}_\theta^{(k)}(\mathbf{x}_t, t),$$

where $\widehat{\mathbf{s}}_\theta^{(k)}(\cdot)$ corresponds to the model trained with the dataset constructed in the same way as ACUR, ω is a hyperparameter controlling the strength of extrapolation.

⁵We downsample each image to 64×64 and define 4 classes based on the intersection of two attributes: “Male” and “Smile.”

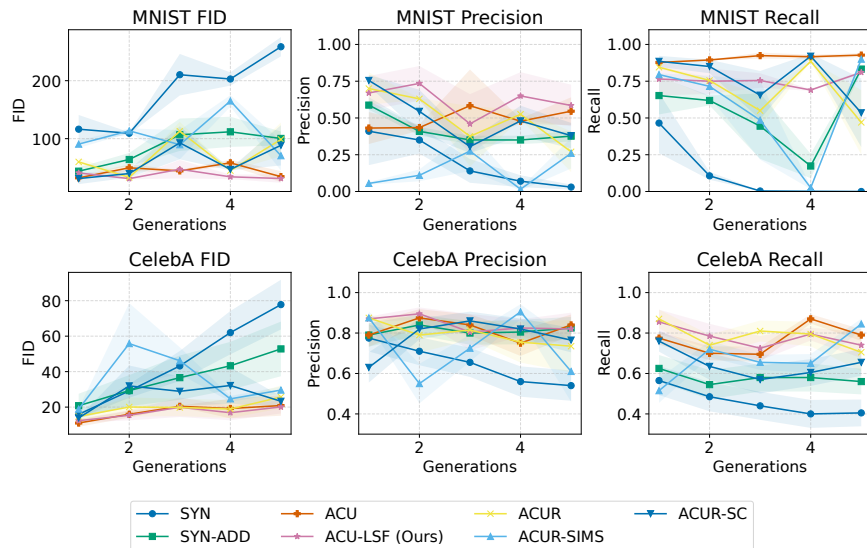


Figure 5: Performance comparison on MNIST and CelebA across three metrics: FID, precision, and recall. FID measures distributional distance between real and synthetic images (lower is better), while precision and recall assess fidelity and diversity of generated samples, respectively (higher is better). SYN suffers from model collapse. SYN-ADD partially alleviates it via fresh real data. ACU, ACUR, and ACU-LSF maintain stable metrics across generations. ACU-LSF achieves lower FID and higher fidelity than ACUR. ACUR-SIMS exhibits instability, while ACUR-SC suffers from reduced recall and elevated FID.

The synthetic dataset $\hat{D}^{(k)}$ is sampled using this new score function $s_{\theta}^{(k)}(\cdot)$.

6. **ACUR-SC** (Gillman et al. 2024): This method clusters real images to obtain class centers and maps each synthetic image to the closest center during training. Besides, the training set is constructed in the same way as ACUR.

Evaluation metrics. Each generation uses 1,000 samples for fine-tuning (except ACU, which accumulates an additional 1,000 images per generation). Training is conducted for three epochs per generation. After each generation, 10,000 images are sampled to evaluate three standard metrics: (i) **Fréchet Inception Distance (FID)** (Heusel et al. 2017) which quantifies the distance between real and synthetic image distributions. (ii) **Precision** (Kynkäänniemi et al. 2019) which measures how many synthetic samples fall on the real data manifold (fidelity). (iii) **Recall** (Kynkäänniemi et al. 2019) which measures how many real samples fall on the synthetic manifold (diversity). Note that lower FID and higher precision and recall values indicate better generative performance. Samples of synthetic images generated by self-consuming models are provided in Fig. 7 (Appendix C) to show the model quality.

Results. We primarily present results for MNIST and CelebA in Figure 5, while additional results for CIFAR-10 are provided in Appendix C. As expected, performance under SYN deteriorates over generations due to model collapse. SYN-ADD partially mitigates collapse by introducing fresh real data. Prior work has shown that increasing the real-to-synthetic ratio improves performance (Kazdan et al. 2024), but collecting fresh real data may not always be feasible. Our goal is to demonstrate that effective filtering without requiring additional real data can offer comparable benefits. ACU maintains stable performance across generations, in

line with previous observations (Gerstgrasser et al. 2024). By randomly sampling from the accumulated dataset, **ACUR** approximates the performance of ACU while adhering to a fixed training budget. As demonstrated in Figure 4, our method **ACU-LSF** successfully removes recent synthetic images and retains more real images, resulting in a distribution that more closely aligns with the original real data. Consequently, ACU-LSF achieves comparable precision and recall, and a lower FID than other fixed-budget baselines. Besides, ACU retrains a diffusion model with 32 M parameters on 582 MB of CelebA data (5 K images at generation 5), and the cost increases with training budget and generations. In contrast, ACU-LSF trains a 65 K-parameter probing classifier on 109 MB of features once at the start, greatly reducing storage and computation.

In contrast, **ACUR-SIMS** and **ACUR-SC** perform less favorably, even when combined with accumulation. SIMS (Ale-mohammad et al. 2024b) was designed to improve fidelity within a single generation via extrapolation. However, over multiple generations, its score functions become unstable due to repeated extrapolation, potentially drifting further from the real data distribution. As a result, metrics fluctuate and performance deteriorates. ACUR-SC employs a self-correction mechanism by aligning synthetic samples with cluster centers of real images. While this exploits structural information from real data, it reduces sample diversity, leading to lower recall and higher FID despite moderate precision.

Overall, our results demonstrate that LSF based on latent representation quality is an effective alternative to adding real data or accumulation, offering a scalable and practical solution to mitigate model collapse in self-consuming training.

Acknowledgements

This work was funded in part by the National Science Foundation under award number IIS-2202699, IIS-2416895, and IIS-2416896.

References

- Alemohammad, S.; Casco-Rodriguez, J.; Luzzi, L.; Humayun, A. I.; Babaei, H.; LeJeune, D.; Siahkoochi, A.; and Baraniuk, R. 2024a. Self-Consuming Generative Models Go MAD. In *The Twelfth International Conference on Learning Representations*.
- Alemohammad, S.; Humayun, A. I.; Agarwal, S.; Colloso, J.; and Baraniuk, R. 2024b. Self-improving diffusion models with synthetic data. *arXiv preprint arXiv:2408.16333*.
- Bertrand, Q.; Bose, J.; Duplessis, A.; Jiralerspong, M.; and Gidel, G. 2024. On the Stability of Iterative Retraining of Generative Models on their own Data. In *The Twelfth International Conference on Learning Representations*.
- Chen, Y.; Liu, J.; Peng, L.; Wu, Y.; Xu, Y.; and Zhang, Z. 2024. Auto-encoding variational Bayes. *Cambridge Explorations in Arts and Sciences*, 2(1).
- Corvi, R.; Cozzolino, D.; Poggi, G.; Nagano, K.; and Verdoliva, L. 2023. Intriguing properties of synthetic images: from generative adversarial networks to diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 973–982.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Feng, Y.; Dohmatob, E.; Yang, P.; Charton, F.; and Kempe, J. 2025. Beyond Model Collapse: Scaling Up with Synthesized Data Requires Verification. In *The Thirteenth International Conference on Learning Representations*.
- Fu, S.; Wang, Y.; Chen, Y.; Tian, X.; and Tao, D. 2025. A Theoretical Perspective: How to Prevent Model Collapse in Self-consuming Training Loops. In *The Thirteenth International Conference on Learning Representations*.
- Fu, S.; Zhang, S.; Wang, Y.; Tian, X.; and Tao, D. 2024. Towards Theoretical Understandings of Self-Consuming Generative Models. In *International Conference on Machine Learning*, 14228–14255. PMLR.
- Gerstgrasser, M.; Schaeffer, R.; Dey, A.; Rafailov, R.; Sleight, H.; Hughes, J.; Korbak, T.; Agrawal, R.; Pai, D.; Gromov, A.; et al. 2024. Is model collapse inevitable? breaking the curse of recursion by accumulating real and synthetic data. *arXiv preprint arXiv:2404.01413*.
- Gillman, N.; Freeman, M.; Aggarwal, D.; Hsu, C.-H.; Luo, C.; Tian, Y.; and Sun, C. 2024. Self-Correcting Self-Consuming Loops for Generative Model Training. In *ICML*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Huang, L.; Qiu, Q.; and Sapiro, G. 2025. SSOLE: Rethinking Orthogonal Low-rank Embedding for Self-Supervised Learning. In *The Thirteenth International Conference on Learning Representations*.
- Kazdan, J.; Schaeffer, R.; Dey, A.; Gerstgrasser, M.; Rafailov, R.; Donoho, D. L.; and Koyejo, S. 2024. Collapse or Thrive? Perils and Promises of Synthetic Data in a Self-Generating World. *arXiv preprint arXiv:2410.16713*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Kwon, M.; Jeong, J.; and Uh, Y. 2023. Diffusion Models Already Have A Semantic Latent Space. In *The Eleventh International Conference on Learning Representations*.
- Kynkäänniemi, T.; Karras, T.; Laine, S.; Lehtinen, J.; and Aila, T. 2019. Improved precision and recall metric for assessing generative models. *Advances in neural information processing systems*, 32.
- LeCun, Y.; Cortes, C.; and Burges, C. 2010. MNIST handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2.
- Lezama, J.; Qiu, Q.; Musé, P.; and Sapiro, G. 2018. Ole: Orthogonal low-rank embedding—a plug and play geometric loss for deep learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8109–8118.
- Li, X.; Zhang, Z.; Li, X.; Chen, S.; Zhu, Z.; Wang, P.; and Qu, Q. 2025. Understanding Representation Dynamics of Diffusion Models via Low-Dimensional Modeling. *arXiv preprint arXiv:2502.05743*.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Lu, Z.; Huang, D.; BAI, L.; Qu, J.; Wu, C.; Liu, X.; and Ouyang, W. 2023. Seeing is not always believing: Benchmarking Human and Model Perception of AI-Generated Images. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Sharifzadeh, S.; Kaplanis, C.; Pathak, S.; Kumaran, D.; Ilic, A.; Mitrovic, J.; Blundell, C.; and Banino, A. 2024. Synth²: Boosting Visual-Language Models with Synthetic Captions and Image Embeddings. *arXiv:2403.07750*.
- Shumailov, I.; Shumaylov, Z.; Zhao, Y.; Gal, Y.; Papernot, N.; and Anderson, R. 2023. The curse of recursion: Training on generated data makes models forget. *arXiv preprint arXiv:2305.17493*.

Shumailov, I.; Shumaylov, Z.; Zhao, Y.; Papernot, N.; Anderson, R.; and Gal, Y. 2024. AI models collapse when trained on recursively generated data. *Nature*, 631(8022): 755–759.

Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32.

Taori, R.; and Hashimoto, T. B. 2023. Data Feedback Loops: Model-Driven Amplification of Dataset Biases. In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*. JMLR.org.

Wang, Z.; Wu, Z.; Zhang, J.; Jain, N.; Guan, X.; and Koshiyama, A. 2024. Bias Amplification: Language Models as Increasingly Biased Media. *arXiv preprint arXiv:2410.15234*.

Wei, X.; and Zhang, X. 2025. Self-Consuming Generative Models with Adversarially Curated Data. In *Forty-second International Conference on Machine Learning*.

Welling, M.; and Teh, Y. W. 2011. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, 681–688. Citeseer.

Wyllie, S.; Shumailov, I.; and Papernot, N. 2024. Fairness Feedback Loops: Training on Synthetic Data Amplifies Bias. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT '24*, 2113–2147. New York, NY, USA: Association for Computing Machinery.

Xie, T.; and Zhang, X. 2024. Automating Data Annotation under Strategic Human Agents: Risks and Potential Solutions. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

You, C.; Robinson, D.; and Vidal, R. 2016. Scalable sparse subspace clustering by orthogonal matching pursuit. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3918–3927.

Yuan, X.; Zhang, C.; Liu, Z.; Shi, D.; Vosoughi, S.; and Lee, W. 2025. Superficial Self-Improved Reasoners Benefit from Model Merging. *arXiv:2503.02103*.