

# Magnifying the Three Phases of GAN Training — Fitting, Refining and Collapsing

Anonymous authors

Paper under double-blind review

## Abstract

Generative Adversarial Networks (GANs) are efficient generative models but may suffer from mode mixture and mode collapse. We present an original global characterization of GAN training by dividing it into three successive phases — fitting, refining, and collapsing. Such a characterization underscores a strong correlation between mode mixture and the refining phase, as well as mode collapse and the collapsing phase. To analyze the causes and features of each phase, we propose a novel theoretical framework that integrates both continuous and discrete aspects of GANs, addressing a gap in existing literature that predominantly focuses on only one aspect. We develop a specialized metric to detect the phase transition from refining to collapsing and integrate it in an “early stopping” algorithm to optimize GAN training. Experiments on synthetic datasets and real-world datasets including MNIST, Fashion MNIST and CIFAR-10 substantiate our theoretical insights and highlight the efficacy of our algorithm.

## 1 Introduction

Generative Adversarial Networks (GANs) serve as a popular technique for unsupervisedly learning generative models of structured and complicated data (Goodfellow et al., 2014). GANs typically involve a generator that produces samples resembling a target dataset, and a discriminator that differentiates between real and generated samples.

One of the main challenges with GAN training is fine-tuning the combined dynamics of the discriminator and generator. Many troublesome phenomena can occur if the dynamics of the generator are not adequately matched with that of the discriminator. Among them, mode mixture (Lei et al., 2019; An et al., 2020; Tanielian et al., 2020) and mode collapse (Goodfellow, 2017) are most commonly observed. Mode collapse occurs when the generator generates limited sample varieties, while mode mixture happens when it blends different modes, producing unrealistic data.

Numerous variants of GAN have been designed to tackle these challenges through a careful blend of hyperparameter optimization and heuristics (Ioffe & Szegedy, 2015; Nowozin et al., 2016; Arjovsky et al., 2017; Li et al., 2017; Nguyen et al., 2017; Ghosh et al., 2018; Miyato et al., 2018). Another line of research focuses on developing theoretical frameworks that can better analyze and optimize GAN training (Lei et al., 2019; An et al., 2020; Sun et al., 2020; Gu et al., 2021; No et al., 2021; Becker et al., 2022; Huang & Zhang, 2023). Regarding theoretical analyses, existing literatures mainly center on the local behavior of GANs near stationary points (Sun et al., 2020; Becker et al., 2022) or the static landscape of GANs (Lei et al., 2019; An et al., 2020; Gu et al., 2021; No et al., 2021). And many of them make straightforward assumptions about the generator (No et al., 2021), the discriminator (No et al., 2021; Becker et al., 2022), or the distribution of real data (Sun et al., 2020; Becker et al., 2022). Please refer to appendix A for additional literature review.

In this paper, we present an original theoretical framework that characterizes the global behavior of GAN dynamics. We divide the training progress into three successive phases called fitting, refining, and collapsing. For each phase, we provide a thorough analysis using well-founded mathematical explanations, backed by elaborate experiments.

## 1.1 Motivation

By training the Non-Saturating GAN (Goodfellow et al., 2014) (NSGAN) on a 3-dimensional Gaussian mixture dataset, we record down the generated samples in fig. 1. The blue dots stand for real samples drawn from the Gaussian mixture, while the generated samples are represented by orange dots. Initially, the generated samples concentrate near the origin. As training progresses, they diffuse, eventually filling the space spanned by real modes. The mode mixture issue is now the most severe. Subsequently, the generated samples are drawn near the modes and the straight lines connecting them. This process refines the generated samples and partially mitigate mode mixture. Next comes the *unexpected collapse*: the generated samples collapse rather than getting more delicate. Every multiple epochs, the number of modes in the generated samples is halved, going from eight to four, then to two, and eventually only one mode remains. This phenomenon is seldom addressed in the existing literature. Nevertheless, it challenges the conventional belief that mode collapse indicates the failure of GAN training. Indeed, if we can terminate GAN training at the opportune moment, the generated samples may showcase considerable diversity.




Figure 1: The real and generated samples by training NSGAN on a 3-dimensional Gaussian mixture dataset. **Orange**: Real samples. **Blue**: Generated samples. Epochs from left to right: 0, 10, 50, 200, 300, 360. Initially, the generated samples cluster near the origin, then spread out and occupy the space spanned by real modes. However, instead of becoming more refined, they eventually collapse to part of the modes.

Through generalization of analogous phenomena across various datasets, we introduce a novel framework delineating the three phases of GAN training: *fitting*, *refining*, and *collapsing*. Fitting corresponds to the first two subfigures in fig. 1. Refining relates to the third subfigure. And collapsing aligns with the last three subfigures. The degree of mode mixture and mode collapse differs across the three phases, as described in table 1. The transition from fitting to refining would refine the generated samples, in terms of reducing both mode mixture and mode collapse. On the contrary, when transiting from refining to collapsing, the quality of generated samples deteriorates — the generator begins to produce delicate yet similar samples, which is undesirable. Therefore, we would like to precisely identify the phase transition from refining to collapsing, and halt training at the right time. Stopping earlier results in more unrealistic samples, whereas stopping later leads to less diverse ones.

Table 1: Severity of mode mixture and mode collapse in three phases. Transitioning from fitting to refining improves sample quality by reducing mode mixture and collapse. Conversely, transitioning from refining to collapsing deteriorates sample quality, resulting in delicate yet similar samples.

	Mode Mixture	Mode Collapse
Fitting	Severe	Severe
Refining	Medium	Medium
Collapsing	Mild	Severe

## 1.2 Our Contributions

Our contributions are threefold:

**A novel theoretical framework** based on NSGAN that integrates discrete and continuous facets of GANs through the use of particle models. Specifically speaking, in section 2.1, we propose a technique called *continuous data augmentation* to augment datasets from discrete samples so that they have continuous

probability density functions. In section 2.2, we discuss about the discrete facet of NSGAN by interpreting it as a *particle model*.

**A comprehensive perspective of GANs** by presenting the novel tripartite phases of GAN training: fitting, refining, and collapsing. We highlight the characteristics of each phase through rigorous mathematical formulations and numerical experiments, as detailed in section 3, section 4, and section 5.

**A specialized metric and the induced “early stopping” algorithm** that optimizes GAN training by detecting the phase transition from refining to collapsing. We elaborate on this metric in section 5.2 and the early stopping algorithm in section 5.3. Notably, this metric is intrinsic to GAN training as it relies on information from the current state of the generator and discriminator, facilitating evaluation and optimization without direct reliance on generated or real images.

## 2 Preparatory Work

This section is dedicated to the preparatory work, including our theoretical framework and problem settings. We first discuss how to model real-world datasets in section 2.1. Therein, we introduce a novel approach called *continuous data augmentation*. We move on to discuss how to interpret Divergence GANs, particularly NSGAN, as a particle model in section 2.2. The fusion of the two methodologies forms our theoretical framework which is well-suited for analyzing GAN dynamics.

### 2.1 How to Model Real-World Datasets?

The probability distributions of real-world datasets are usually modeled as a linear combination of Dirac measures that remain unchanged throughout GAN training (Sun et al., 2020; Becker et al., 2022). This kind of modeling may not be optimal since neural networks are typically trained from batches of data, and the batch size is relatively small compared with the size of the dataset. Therefore, the distribution of data varies substantially from batch to batch. To address such a misalignment, we propose *continuous data augmentation*, a way to augment datasets so that they have continuous probability density functions.

**Definition 2.1.** Let  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N \in \mathbb{R}^n$ . Given a discrete probability measure  $\mu = \frac{1}{N} \sum_{i=1}^N \delta_{\mathbf{x}_i}$ , where  $\delta_{\mathbf{x}_i}$  is the Dirac measure with unit mass concentrated on the point  $\mathbf{x}_i$ . The continuous data augmentation of  $\mu$  is defined as a continuous probability measure  $\hat{\mu}$  with density function

$$\hat{f}(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N K_h(\mathbf{x}, \mathbf{x}_i)$$

where  $K_h(\mathbf{x}, \mathbf{y}) = K(\mathbf{x}/h, \mathbf{y}/h)/h$  is the scaled kernel with bandwidth  $h > 0$ .

We always assume, if not otherwise stated, that  $K$  is the Radial Basis Function (RBF) kernel

$$K(\mathbf{x}, \mathbf{y}) = \left( \sqrt{1/\pi} \right)^n \exp(-\|\mathbf{x} - \mathbf{y}\|_2^2).$$

This kind of formulation approximates the distribution of real-world datasets in a flexible way. When the bandwidth  $h$  tends to 0,  $\hat{\mu}$  converges to the original probability measure  $\mu$ . When  $h$  is away from 0, samples from the continuously augmented dataset are more diversified while remaining close to certain real samples from the original dataset. Moreover, when training GANs, batches of data can be viewed as random samples from the augmented dataset, rather than remaining fixed throughout training, which better mirrors reality.

### 2.2 Rethinking Divergence GANs as Particle Models

Divergence GANs such as Vanilla GAN (Goodfellow et al., 2014), NSGAN (Goodfellow et al., 2014) and  $f$ -GAN (Nowozin et al., 2016) can be interpreted as *particle models* (Gao et al., 2019; Franceschi et al., 2023; Huang & Zhang, 2023). This paper focuses on the NSGAN for its practicality and conciseness. And we outline the methodology for other Divergence GANs in appendix H. The pseudocode of NSGAN as a

**Algorithm 1** Interpretation of Non-Saturating GAN (NSGAN) as a Particle Model**Require:** The discriminator  $d_\omega$  and the generator  $g_\theta$ , the noise prior  $p_z$ , batch size  $m > 0$ , stepsize  $s > 0$ 

- 1: **for** number of training iterations **do**
- 2:   Train the discriminator  $d_\omega$  as in (Goodfellow et al., 2014).
- 3:   Sample  $\mathbf{z}_i$ 's from the noise prior  $p_z(\mathbf{z})$ .
- 4:   Generate particles

$$\mathbf{Z}_i = g_\theta(\mathbf{z}_i), \quad i = 1, \dots, m.$$

- 5:   Update the particles

$$\hat{\mathbf{Z}}_i = \mathbf{Z}_i + s \cdot \frac{\nabla d_\omega(\mathbf{Z}_i)}{2d_\omega(\mathbf{Z}_i)}, \quad i = 1, \dots, m.$$

- 6:   Update  $g_\theta$  by descending its stochastic gradient:

$$\nabla_\theta \frac{1}{m} \sum_{i=1}^m \|g_\theta(\mathbf{z}_i) - \hat{\mathbf{Z}}_i\|_2^2.$$

- 7: **end for**

particle model is presented in algorithm 1. We show in theorem 2.1 (with the proof deferred to appendix C) that algorithm 1 is essentially equivalent to the original NSGAN. Thus we also refer to generated samples as *particles* in this paper.

**Theorem 2.1.** *The update of  $g_\theta$  via descending the gradient*

$$\nabla_\theta \frac{1}{m} \sum_{i=1}^m \|g_\theta(\mathbf{z}_i) - \hat{\mathbf{Z}}_i\|_2^2$$

*in algorithm 1 is equivalent to descending the gradient*

$$-\nabla_\theta \frac{1}{m} \sum_{i=1}^m \log(d_\omega(g_\theta(\mathbf{z}_i)))$$

*in the original formulation of NSGAN.*

We make the following assumption on the noise prior  $p_z(\mathbf{z})$  for reasons in appendix B.

**Assumption 2.1.** *Let  $n$  be the dimension of real samples. We assume that the noise prior  $p_z \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$  is an  $n$ -dimensional standard Gaussian distribution.*

### 3 The First Phase of GAN Training — Fitting

We begin with the first phase of GAN training: fitting. Fitting refers to the process where the generated samples progressively spread to cover the space that envelopes the modes. We first analyze how particles are updated in section 3.1. Therein, we identify two regimes where particles exhibit markedly different behaviors during updating. Next, we employ a multiscale approach to model real-world distributions by building two models at different scales. Ensuring consistency across the analysis, the 3-dimensional Gaussian mixture used in section 1.1 serves as the basis. We introduce two minor adjustments to the original setup. Firstly, the 3-dimensional distribution is projected onto the  $xy$ -plane and its associated marginal distribution is used for clearer visualization. Secondly, the covariance matrix of each Gaussian component is adjusted to  $0.1\mathbf{I}$  to amplify the effect. The two models are elucidated respectively in section 3.2 and section 3.3. Throughout this section, we assume that the discriminator is optimal, i.e.,

$$d_\omega^*(\mathbf{x}) = \frac{p_{\text{data}}(\mathbf{x})}{p_{\text{data}}(\mathbf{x}) + p_g(\mathbf{x})}.$$

as established by Goodfellow et al. (2014).






Figure 2: The vector field  $\nabla d_\omega^*(\mathbf{x})/d_\omega^*(\mathbf{x})$  and the values of the discriminator in different models. **First:** Model 1 at initialization. **Second:** Model 1 with generated samples covering all modes. **Third:** Model 1 with generated samples covering only one mode. **Fourth:** Model 2.

### 3.1 How Particles are Updated

Recall in algorithm 1 that a particle  $\mathbf{x}$  is updated following

$$\mathbf{x} = \mathbf{x} + s \cdot \frac{\nabla d_\omega(\mathbf{x})}{2d_\omega(\mathbf{x})},$$

where  $s$  is the stepsize. We first theoretically explore the properties of  $\nabla d_\omega(\mathbf{x})/d_\omega(\mathbf{x})$  in theorem 3.1, deferring its proof to appendix C. In this theorem, we derive two regimes corresponding to the limiting scenarios of  $p_{\text{data}}(\mathbf{x})/p_g(\mathbf{x})$ , which lead to significantly different update mechanisms for the particles.

**Theorem 3.1.** *Assume that the discriminator is optimal, i.e.,  $d_\omega^*(\mathbf{x}) = p_{\text{data}}(\mathbf{x})/(p_{\text{data}}(\mathbf{x}) + p_g(\mathbf{x}))$ . Denote  $r(\mathbf{x}) = p_{\text{data}}(\mathbf{x})/p_g(\mathbf{x})$ . At a point  $\mathbf{x}$  where  $r(\mathbf{x}) \approx 0$ ,  $\mathbf{x}$  is updated following approximately  $\nabla \log(r(\mathbf{x}))$ . Conversely, when  $r(\mathbf{x}) \gg 1$ ,  $\mathbf{x}$  is updated following approximately  $\nabla(-1/r(\mathbf{x}))$ .*

We hereby outline the implications of theorem 3.1. The value of  $\log(r(\mathbf{x}))$  changes dramatically as  $\mathbf{x}$  decreases from 1 to 0, leading to correspondingly large magnitudes of  $\|\nabla \log(r(\mathbf{x}))\|_2$  when  $r(\mathbf{x}) \approx 0$ . This indicates that in the regions where  $p_g(\mathbf{x})$  significantly exceeds  $p_{\text{data}}(\mathbf{x})$ , particles are propelled towards distant points. Conversely,  $\nabla(-1/r(\mathbf{x}))$  changes more gradually with increasing  $\mathbf{x}$ , resulting in smaller magnitudes of  $\|\nabla(-1/r(\mathbf{x}))\|_2$  when  $r(\mathbf{x}) \gg 1$ . In such regions where  $p_g(\mathbf{x})$  is lower than  $p_{\text{data}}(\mathbf{x})$ , particles tend to remain relatively stationary. Below we will see how empirical observations align with the theoretical results.

### 3.2 Model 1: The Modes are Clustered

We first examine the scenario where the modes are clustered. Such a configuration is common in real-world datasets in the *local* sense. For instance, the two handwritten digits 1 and 7 are similar and their modes are close in MNIST (LeCun et al., 1998) (please refer to appendix D for a visualization of how close they are). In this subsection, we investigate three typical intermediate stages of GAN training: initialization, where generated samples cover all the modes, and where generated samples cover only one mode.

For initialization, we assume that  $p_g$  equals the Gaussian distribution  $\mathcal{N}([0, 0], 0.2\mathbf{I}_2)$  and

$$p_{\text{data}} \sim \frac{1}{4}\mathcal{N}([1, 1], 0.1\mathbf{I}_2) + \frac{1}{4}\mathcal{N}([1, -1], 0.1\mathbf{I}_2) + \frac{1}{4}\mathcal{N}([-1, 1], 0.1\mathbf{I}_2) + \frac{1}{4}\mathcal{N}([-1, -1], 0.1\mathbf{I}_2).$$

The vector field that a particle is updated (i.e.,  $\nabla d_\omega^*(\mathbf{x})/d_\omega^*(\mathbf{x})$ ) is plotted in the first subfigure of fig. 2. We observe that the particles are drawn towards the nearest modes, and there is a positive correlation between the vector lengths, i.e., the intensity of attraction and the particles' distances from the modes.

For the second case where the generated samples cover all the modes, we modify  $p_g$  to follow the uniform distribution on  $[-2, 2] \times [-2, 2]$ . The vector field that a particle is updated (i.e.,  $\nabla d_\omega^*(\mathbf{x})/d_\omega^*(\mathbf{x})$ ) is plotted in the second subfigure of fig. 2. We observe that the particles close to the centers of the modes tend to remain stationary, whereas particles far from the modes will be updated to bring them closer to the nearest mode, which aligns with our discussion in section 3.1.

Finally, we investigate the case where generated samples cover a single mode. Assume that  $p_g \sim \mathcal{N}([1, 1], \mathbf{I}_2)$ , so that the generated samples cover the mode centered at  $(1, 1)$ . Please refer to the third subfigure of fig. 2 for the plot of  $\nabla d_\omega^*(\mathbf{x})/d_\omega^*(\mathbf{x})$ . This time, the discriminator values are unevenly distributed near the four modes: the lowest for the mode covered by generated samples, and the highest for the mode furthest away from the generated samples. Moreover, the intensity of attraction, as indicated by vector lengths, peaks near the mode unoccupied by generated samples but diminishes near the crowded one. The intuition is that once a particle come into proximity to the unoccupied mode, it will be forcefully drawn towards it.

### 3.3 Model 2: The Modes are Far Apart

In this subsection, we study the case where the modes are far apart. From a *global* perspective, real-world datasets often exhibit such a structure, particularly those with multiple categories. As an illustration, the modes of dogs and trucks in CIFAR-10 (Krizhevsky et al., 2009) are far apart because of the significant differences in their appearances (please refer to appendix D for a visualization of how far they are).

We demonstrate the difficulty for the generator to capture all the modes in such a case. Assume that

$$p_{\text{data}} \sim \frac{1}{4}\mathcal{N}([3, 3], 0.1\mathbf{I}_2) + \frac{1}{4}\mathcal{N}([3, -3], 0.1\mathbf{I}_2) + \frac{1}{4}\mathcal{N}([-3, 3], 0.1\mathbf{I}_2) + \frac{1}{4}\mathcal{N}([-3, -3], 0.1\mathbf{I}_2)$$

and  $p_g \sim \mathcal{N}([3, 3], 3\mathbf{I}_2)$ . The plot of  $\nabla d_\omega^*(\mathbf{x})/d_\omega^*(\mathbf{x})$  can be found in the last subfigure of section 3.3. We notice a general weakening of attraction intensity near all modes, presenting challenges for particles to move towards an unoccupied mode subject to the vector field. This situation could result in a “pre-refining” mode collapse or nonconvergence, preventing GAN training from advancing to the refining phase. Such collapse or nonconvergence might be attributed to detrimental network initialization or imbalanced training of the generator and the discriminator. Given its infrequency in practice, we will not delve deeper into this topic.

## 4 The Second Phase of GAN Training — Refining

This section discusses the refining phase of GAN training where the generated samples become more refined. We delve into the relationship between sample quality and the “steepness” of the generator function, a concept that we will define later. We demonstrate that to push  $p_z$  towards a multimodal distribution  $p_{\text{data}}$ , the generator function  $g$  must exhibit significant steepness in section 4.1. While the steepness does evolve over the course of training, it does so at a slow pace. This gradual evolution contributes to the refinement of generated samples but also allows mode mixture to remain a prevalent issue, as we explore in section 4.2.

### 4.1 Steepness of Measure-Preserving Maps

Following the *continuous data augmentation*, we assume in this subsection that  $p_{\text{data}}$  is a 1-dimension Gaussian mixture distribution, i.e.,

$$p_{\text{data}} \sim \alpha_1\mathcal{N}(\mu_1, \sigma^2) + \cdots + \alpha_n\mathcal{N}(\mu_n, \sigma^2).$$

We present a characterization of the measure-preserving map  $g$  that satisfies  $g_\#p_z = p_{\text{data}}$  in theorem 4.1. The proof is deferred to appendix C.

**Theorem 4.1.** *Let  $\Phi(x)$  denotes the cumulative distribution function (CDF) of  $\mathcal{N}(0, 1)$  and let  $\Psi(x)$  be that of  $p_{\text{data}}(x)$ . If  $g$  satisfies  $g_\#p_z = p_{\text{data}}$ , then  $g = \Psi^{-1} \circ h \circ \Phi$ , where  $h$  is a measure-preserving map of  $\mathcal{U}(0, 1)$ , i.e., the uniform distribution on  $(0, 1)$ .*

Among the many measure-preserving maps with different  $h$ ’s, special attention is given to where  $h$  equals the identity map. In this instance, the corresponding  $g$  is the optimal transport from  $p_z$  to  $p_{\text{data}}$  under the 2-Wasserstein distance (Santambrogio, 2015). We start examining the properties of  $g = \Psi^{-1} \circ \Phi$  by presenting a comprehensive visualization of  $g$  given different  $p_{\text{data}}$  in fig. 3. Notably, several techniques are employed in crafting these figures and details can be found in appendix I. We notice that there is a peak in the magnitude of the derivative  $g'$  between each pair of adjacent modes. This observation is explicitly articulated in theorem 4.2, wherein we consider a simplified scenario involving only two Gaussians in  $p_{\text{data}}$ . Again, the proof is deferred to appendix C.




Figure 3: The Functions  $g$  that satisfy  $g_{\#}p_z = p_{\text{data}}$ , where  $p_{\text{data}}$  is a Gaussian mixture. **First:** The mean  $\mu$  affects  $g'(0)$ . **Second:** The variance  $\sigma^2$  affects the asymptotic slope of  $g$  away from zero. **Third:** The mixing coefficients  $\{\alpha_i\}_{i=1}^n$  affect the inclination of  $g$  towards specific modes. **Fourth:** The number of Gaussians  $n$  affects the number of segments in  $g$ .

**Theorem 4.2.** Assume that

$$p_{\text{data}} \sim 0.5\mathcal{N}(-\mu, 1) + 0.5\mathcal{N}(\mu, 1).$$

Then  $g(x) := \Psi^{-1}(\Phi(x))$  satisfies  $g'(0) = \exp(\mu^2/2)$ .

We conclude that the magnitude of  $g'$  depends on the distance between two modes in 1-dimensional cases. Before exploring high-dimensional scenarios, we introduce the concept of “steepness” of a function that generalizes the notion of derivative in the 1-dimensional case.

**Definition 4.1.** Let  $g : \mathbb{R}^n \rightarrow \mathbb{R}^n$  be continuously differentiable. We define its steepness  $\mathcal{S}(g)$  as the upper bound of the max norm of its Jacobian, namely,

$$\mathcal{S}(g) = \sup_{\mathbf{x} \in \mathbb{R}^n} \|J_g(\mathbf{x})\|_{\max}.$$

We are now in a position to demonstrate the universality of this phenomenon in higher dimensions. Specifically speaking, for a function  $g$  to satisfy  $g_{\#}p_z = p_{\text{data}}$  where  $p_z$  is standard Gaussian and  $p_{\text{data}}$  is a distribution with Radial Basis Function kernel density (as in section 2.1), we prove in theorem 4.3 that  $\mathcal{S}(g)$  has a lower bound which depends on the dimension  $n$ , the bandwidth  $h$ , and the  $\mathbf{x}_i$ ’s. The proof is rather technical and complicated, and can be found in appendix C.

**Theorem 4.3.** Let  $\nu$  be the truncated Gaussian distribution  $\mathcal{N}_r(\mathbf{0}, \mathbf{I}_n)$  in the  $n$ -dimensional ball  $\mathcal{B}_r(\mathbf{0})$  and assume that  $\hat{\mu}$  is a probability measure with probability density function

$$\hat{f}(\mathbf{x}) = \frac{1}{Nh} \sum_{i=1}^N \left( \sqrt{1/\pi} \right)^n \exp \left( -\frac{\|\mathbf{x} - \mathbf{x}_i\|_2^2}{h^2} \right).$$

Suppose that  $g : \mathcal{B}_r(\mathbf{0}) \rightarrow \mathbb{R}^n$  is continuously differentiable and piecewise injective. Then  $\mathcal{S}(g) > M$ , where

$$M = \delta \cdot h^{1/n} \cdot \sqrt{\frac{\pi}{n}} \cdot \max_{1 \leq i \leq N} \exp \left( \frac{\|\bar{\mathbf{x}} - \mathbf{x}_i\|_2^2}{nh^2} \right).$$

Here,  $\bar{\mathbf{x}} = \sum_{i=1}^N \mathbf{x}_i / n$ , and  $\delta = \exp(-r^2/2) / \sqrt{2\pi}$ .

Theorem 4.3 gives an explicit lower bound of the steepness. Notably, this lower bound exhibits an exponential dependence on both the Euclidean distance  $\|\bar{\mathbf{x}} - \mathbf{x}_i\|_2$  and the reciprocal of the bandwidth  $1/h$ . When the modes  $\mathbf{x}_i$ ’s are considerably distant from each other or when the bandwidth  $h$  diminishes, the steepness  $\mathcal{S}(g)$  will be large. Consequently, using a neural network to approximate  $g$  poses great challenges, which we will elaborate on in section 4.2.

## 4.2 Evolution of Generator’s Steepness

We use a synthetic example to illustrate how the steepness of the generator evolves during training. Theorem 4.4 gives the recurrent formula in the setting where the optimal generator has a sufficiently large

steepness  $k_*$ . The proof involves examining the trajectory of a designated particle throughout the training process, with detailed verification deferred to appendix C.

**Theorem 4.4.** *Assume that  $p_{data} \sim \mathcal{N}(\mathbf{0}, k_*^2 \mathbf{I}_n)$  and that the discriminator is optimal, i.e., the discriminator consistently provides the precise moving direction for the particle. Then  $k_t$ , the steepness of  $g$  at discrete time step  $t$  satisfies*

$$k_{t+1} = k_t + s \left( \frac{1}{k_t^2} - \frac{1}{k_*^2} \right) \cdot \frac{1}{1 + \frac{k_t \varphi(k_t \mathbf{x}_0 / k_*)}{k_* \varphi(\mathbf{x}_0)}},$$

where  $0 \leq t \leq T$ , and  $T$  is the maximum time. Here,  $\varphi$  is the probability density function of  $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ .

We now explain how refining happens. Following the identical configuration as in theorem 4.2 with  $\mu = 8$ , we model the generator function  $g_k$  as

$$g_k(x) = \begin{cases} x - 8, & x < -8/(k-1), \\ kx, & -8/(k-1) \leq x \leq 8/(k-1), \\ x + 8, & x > 8/(k-1). \end{cases}$$

Here, we use  $k$  to indicate the steepness of  $g$ . The  $x$ -axis can be divided into five disjoint intervals separated by  $-8/(k-1)$ ,  $-4/(k-1)$ ,  $4/(k-1)$  and  $8/(k-1)$ . When  $z \sim p_z$  falls into the interval  $[-4/(k-1), 4/(k-1)]$ , it follows that  $g_k(z)$  lies between the two modes  $-8$  and  $8$  as long as  $k > 3/2$ . This would cause mode mixture. To quantify the severity, we compute the probability that  $z \in [-4/(k-1), 4/(k-1)]$  given  $z \sim p_z$ . Here, we set  $k_* = \exp(32)$ ,  $k_0 = 1$  and  $s = 5$ . We then plot  $k_t$  against  $t$  and the probability against  $k$  respectively in fig. 4. As  $t$  grows, the steepness  $k_t$  initially rises quickly and then slows down. With a higher  $k$ , fewer samples fall between the two modes, thereby facilitating the refining phase. Note that when  $k_t$  hits 250, only about 1.2 percent of the samples are between the two modes.




Figure 4: The evolution of the steepness  $k_t$  against  $t$  and the probability that generated samples fall between the modes as  $k$  progresses. **Left:**  $k_t$  against  $t$ . **Right:** Probability against  $k$ . The steepness  $k_t$  rises quickly before slows down. A higher  $k$  results in fewer samples between the modes, facilitating the refining phase.

## 5 The Third Phase of GAN Training — Collapsing

This section focuses on the third phase of GAN training: collapsing. Rather than improving, the generated samples would eventually collapse to a limited number of modes. We start by offering an empirical understanding of this phenomenon in section 5.1. Following that, we introduce a metric that can detect the phase transition from refining to collapsing in section 5.2. After theoretically deriving such a metric, we propose a novel early stopping algorithm that can judiciously halt GAN training when it nears or reaches the brink of collapse, preventing further deterioration in sample quality in section 5.3.

### 5.1 What Is the Reason for GANs’ Collapsing?

We empirically demonstrate that collapsing stems from the “overfitting” of the discriminator. Overfitting generally describes the situation where a machine learning model makes precise predictions for the training data but fails to do so for unseen data. We employ this idea to denote the discriminator’s propensity to give very high values to real samples and exceedingly low values to other samples. Please refer to appendix F

for a visualization of the optimal discriminator’s such behavior. To see the connection, we focus on the way generated particles are updated in algorithm 1:

$$\hat{\mathbf{Z}} = \mathbf{Z} + s \cdot \frac{\nabla d(\mathbf{Z})}{2d(\mathbf{Z})}.$$

Here,  $d$  represents the discriminator function (omitting the subscript  $\omega$  for brevity),  $\mathbf{Z}$  is a particle before update, and  $\hat{\mathbf{Z}}$  is the same particle after update. If the discriminator “overfits”, there will be a large disparity between its value within the modes and outside of the modes. Consequently, the norm of the gradient  $\|\nabla d\|_2$  will reach its peak at the boundaries of the modes. This will displace  $\hat{\mathbf{Z}}$  significantly from  $\mathbf{Z}$  if  $\mathbf{Z}$  is close to the boundary. The boundaries of the modes will therefore shrink, and this process will persist until all the generated samples have dispersed. Please refer to fig. 5 for an illustration and experimental observations.



Figure 5: A depiction of why collapsing happens. **Left:** An illustration. **Right:** Experimental observations on the 3-dimensional Gaussian mixture dataset. In the illustration, the original particles (orange rhombi) near the boundary of the mode (solid gray circle / sphere) will be pushed to a significant distance (orange-red pentagons). The particles within the mode (green rhombi) will stay close to their original places (dark green pentagons). This contracts the boundary of the mode (gray dotted line).

## 5.2 Detecting the Phase Transition From Refining to Collapsing

In this section, we give a priori estimation of  $\|\nabla d(\mathbf{x})/d(\mathbf{x})\|_2$  in the collapsing phase. Since collapsing manifests at the end of refining, it is important to analyze the discriminator’s behavior when generated and real samples are close. Unlike the optimal discriminator in Vanilla GAN that assigns a uniform value of 0.5 to both real and generated samples when fully trained, the optimal discriminator in NSGAN assigns values near 0.5 to the central regions of modes and values near 0 to regions with scarce real samples. Amidst them, the values gradually diminish from 0.5 to 0. Please refer to appendix F for a visualization of the optimal discriminator’s such behavior. This observation leads to a linear model of the discriminator in assumption 5.1. The rationale behind selecting the radius as  $2\sqrt{2}h$  is given in appendix G.

**Assumption 5.1.** *Let the probability density function of the real distribution be*

$$p_{data}(\mathbf{x}) = \frac{1}{Nh} \sum_{i=1}^N \left( \sqrt{1/\pi} \right)^n \exp \left( -\frac{\|\mathbf{x} - \mathbf{x}_i\|_2^2}{h^2} \right),$$

where we require  $\min_{1 \leq i, j \leq N} \|\mathbf{x}_i - \mathbf{x}_j\|_2 \geq 2\sqrt{2}h$ . We assume that at the end of the refining phase where generated samples closely resemble real samples, the discriminator  $d(\mathbf{x})$  is of the form

$$d(\mathbf{x}) = \begin{cases} \frac{1}{2} - \frac{\sqrt{2}}{8h} \|\mathbf{x} - \mathbf{x}_i\|_2, & \mathbf{x} \in B_{2\sqrt{2}h}(\mathbf{x}_i), \\ 0, & \text{otherwise.} \end{cases}$$

Now, we estimate  $\|\nabla d(\mathbf{x})/d(\mathbf{x})\|_2$  based on assumption 5.1. Considering that collapsing typically takes place near the boundaries of the modes (refer to section 5.1), we calculate  $\|\nabla d(\mathbf{x})/d(\mathbf{x})\|_2$  at  $\tilde{\mathbf{x}}$  that locates  $\sqrt{2}h$  away from  $\mathbf{x}$ . This choice originates from the fact that for Gaussian distributions, 95 percent of the samples fall within a sphere with a radius of two standard deviations from the mean. Through direct computation, we derive that  $\|\nabla d(\tilde{\mathbf{x}})/d(\tilde{\mathbf{x}})\|_2 = \sqrt{2}/(2h)$ . Note that this value solely relies on the bandwidth  $h$  and remains unaffected by the dimension of the dataset.

**Algorithm 2** Early Stopping of GANs

---

**Require:** A GAN model including a generator  $g_\theta$  and a discriminator  $d_\omega$ , the estimated bandwidth  $h > 0$ , the scale factor  $k_s > 0$ , the number of modes  $m \geq 1$ , the number of warm-up training iterations  $N_w$

- 1: **for** number of training iterations **do**
- 2:   Train the discriminator  $d_\omega$  as in algorithm 1.
- 3:   Compute the  $(1 - 1/m)$  quantile of  $\|\nabla d_\omega/d_\omega\|_2$  in a batch. Let the value be  $q$ .
- 4:   **if**  $q > k_s \cdot \sqrt{2}/(2h)$ , current iteration  $> N_w$  **then**
- 5:     **break**
- 6:   **end if**
- 7:   Train the generator  $g_\theta$  as in algorithm 1.
- 8: **end for**

---

### 5.3 Early Stopping

We have theoretically derived  $\|\nabla d(\mathbf{x})/d(\mathbf{x})\|_2$  in the collapsing phase. In this section, we propose a novel algorithm called “early stopping” that can halt GAN training at an early time when it nears or reaches the brink of collapse, preventing further deterioration in sample quality. The pseudocode is presented in algorithm 2.

There are three key ingredients:  $k_s$ ,  $m$  and  $N_w$ . The scale factor  $k_s$  will be multiplied by  $\sqrt{2}/(2h)$  to establish a threshold. Training is terminated if the value of  $\|\nabla d_\omega/d_\omega\|_2$  exceeds this threshold. Then the best-performing model will be selected among the checkpoints saved before the stopping point. The scale factor is set proportional to the distances between two adjacent modes in the dataset. The underlying rationale is that when  $\|\nabla d_\omega/d_\omega\|_2$  is small compared with the inter-mode distances, generated samples that are deviated from the modes can be re-attracted to the modes, causing no collapsing. Nevertheless, as  $\|\nabla d_\omega/d_\omega\|_2$  grows to be comparable to the inter-mode distances, it is increasingly likely for the generated samples to gravitate towards alternate modes, posing a potential risk of mode collapse. The parameter  $m$  denotes the number of modes and serves as a criterion for determining the proportion of samples within a batch that will be compared against the aforementioned threshold. Specifically, we opt to evaluate samples corresponding to the  $(1 - 1/m)$ th quantile. Underlying this choice is the presumption that once a particular mode initiates collapsing, it signifies the commencement of the GAN training transitioning into the collapsing phase. The number of warm-up training iterations  $N_w$  is user-defined. It indicates how many training iterations GANs would undergo before transitioning into the second phase. This provision serves as a safeguard against premature cessation of training because during the fitting phase,  $\|\nabla d_\omega/d_\omega\|_2$  may exceed the threshold as well.

## 6 Experiments

In this section, we present the experimental results. All the codes are available in the supplementary material.

### 6.1 Verifying Fitting and Refining

We demonstrate the existence of fitting and refining in real-world datasets. Our experiments focus on MNIST and Fashion MNIST because of the clarity of their modes. Detailed results as well as results for Fashion MNIST are deferred to appendix F. The experimental settings and rationale can be found in appendix E.

**Methodology.** We train NSGAN on MNIST and use a classification network  $q(\mathbf{x})$  to analyze the generated images. Here,  $\mathbf{x}$  is an image tensor, and  $q(\mathbf{x})$  outputs a 10-dimensional vector  $(p_0, p_1, \dots, p_9)$ , where  $p_i \in [0, 1]$  represents the likelihood of  $\mathbf{x}$  corresponding to the handwritten digit  $i$ . We count the pairs  $(i, j)$  within a batch where both  $p_i$  and  $p_j$  exceed  $10^{-2}$ , and visualize the occurrences using heatmaps in fig. 6. We infer that such  $\mathbf{x}$  exhibits a mixture of modes  $i$  and  $j$  by pairing  $(i, j)$ .

**Results.** Initially, only few entries in the heatmap are nonzero, meaning that the generated images all look similar. As training progresses, more entries become nonzero, reflecting the fitting phase where the generated

samples spread to cover the space that envelopes the modes. The off-diagonal entries serve as indicators of mode mixture. We observe that they decrease in magnitude as training advances, validating the refining phase. However, the issue of mode mixture persists even at the closure of refining. These observations are in line with what we have derived in section 3 and section 4.




Figure 6: The logarithm of the occurrence of pairings  $(i, j)$  plus 1 in a batch of size 256. Epochs from left to right: 0, 1, 2, 4, 8. At first, there are few nonzero entries. As training continues, more entries become nonzero, indicating the spread of generated samples across mode space, i.e., the fitting phase. Off-diagonal entries signal mode mixture, decreasing over training and validating the refining phase. However, mode mixture persists even after refining. The annotated heatmaps can be found in appendix F.

## 6.2 Early Stopping

We present the results of early stopping (algorithm 2) applied to both synthetic and real datasets, including 3-dimensional Gaussian mixture, MNIST, Fashion MNIST and CIFAR-10. The detailed experimental settings are provided in appendix E and the results of more training runs can be found in appendix F.

**Early stopping.** We train NSGAN on different datasets and record down  $\|\nabla d_\omega/d_\omega\|_2$  each epoch until the maximum specified epochs. The threshold in early stopping is determined by  $k_s \cdot \sqrt{2}/(2h)$ , where  $k_s$  equals the estimated distance between two modes and  $h$  equals the estimated bandwidth. Instead of halting training when this norm exceeds the threshold, we opt to continue training, which allows us to assess the sample quality both before and after the stop. The experimental results for the 3-dimensional Gaussian mixture, MNIST, Fashion MNIST, and CIFAR-10 are shown in fig. 8.

**Comparison with choosing the best FID model.** Our early-stopping algorithm complements the common practice to save model checkpoints and choose the best-performing one via the FID score (Heusel et al., 2017). While existing methods use empirical approaches to decide the maximum training epoch, our algorithm guides practitioners on when to halt training, minimizing the need for saved model checkpoints and conserving computational resources. Remarkably, we demonstrate in fig. 9 that our intrinsic metric correlates with the FID score by accurately capturing the surge in FID score indicative of mode collapse.

**Validating the early stopping metric.** We validate the effectiveness of our early stopping metric by demonstrating that, upon applying the technique of injecting noise into the intermediate layers of the discriminator to combat mode collapse, the metric is pushed back. The results can be found in fig. 7.




Figure 7: Histograms of the values of  $\|\nabla d_\omega/d_\omega\|_2$  and their 90th percentile across epochs. Red for the model with noise and blue for the model without noise. The noise-free GAN collapses at the 54th epoch. Preceding the 54th, the noised model nearly always exhibits lower  $\|\nabla d_\omega/d_\omega\|_2$  values compared to its noise-free counterpart. Post 54th epoch, this relationship reverses. Notably, in the noise-free model,  $\|\nabla d_\omega/d_\omega\|_2$  tends towards zero, contributing to this observed divergence. Please refer to appendix F for additional results.



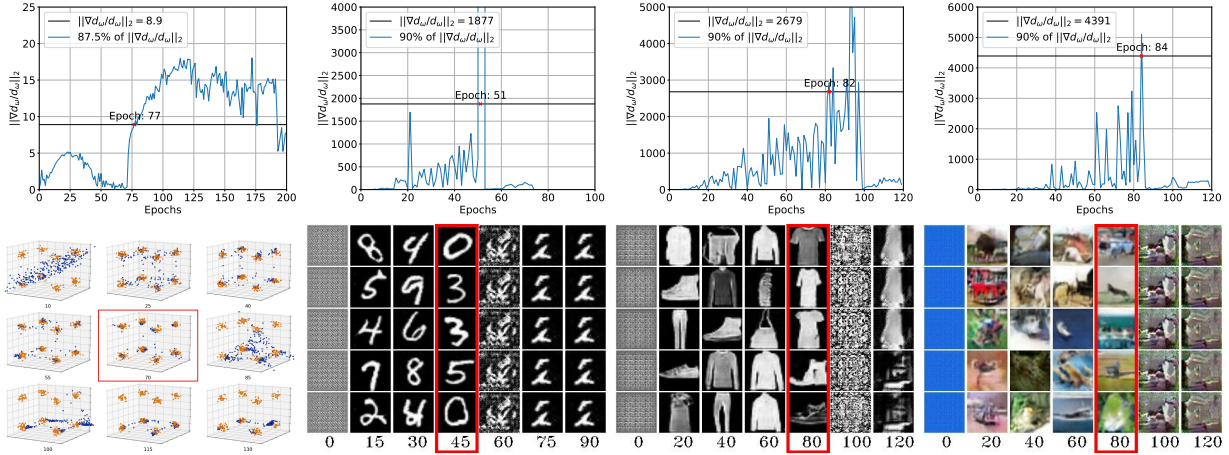


Figure 8: Experimental results of early stopping. The figures in each column, from left to right, represent the results from the following datasets: Gaussian mixture, MNIST, Fashion MNIST, and CIFAR-10, respectively. The images framed in **red** are the most realistic ones within the presented ones. **Gaussian mixture**: The generated samples consistently demonstrate high quality after the fitting phase. As training approaches the stopping point, a discernible deviation in sample quality is observed. Eventually, the samples collapse to half of the modes. **MNIST**: Prior to the stopping point, the generated samples exhibit high quality and variety. After the stopping point, they are first contaminated by noise. And then collapse to a specific mode. **Fashion MNIST**: Before the stopping point, the generated images demonstrate both clarity and diversity. After the stopping point, they first become noise, and then collapse to some of the modes. **CIFAR-10**: Before the stopping point, the generated samples were crisp and varied. After the stopping point, they collapse and transition between different modes. The results of more training runs can be found in appendix F




Figure 9: The tendency of  $\|\nabla d_\omega/d_\omega\|_2$  and the FID scores on MNIST, Fashion MNIST and CIFAR-10, from left to right. **Red circled** for the FID scores and **blue** for our metric. A consistent pattern is observed: Whenever our metric surges past the threshold, the FID scores nearly concurrently escalate to high values, signifying a notable deterioration in sample quality. Please refer to appendix F for additional results.

## 7 Conclusion

This study presents a tripartite framework for GAN training, encompassing the phases of fitting, refining, and collapsing. Through rigorous mathematical formulations and corroborative numerical experiments, we highlight the characteristics of each phase. Our comprehensive examination unveils the intricate dynamics associated with the challenges of mode collapse and mode mixture. Moreover, to enhance the efficiency and robustness of GAN training, we introduced a novel early stopping algorithm that is intrinsic to GAN training and does not rely on extrinsic metrics like the FID score. In the future, we plan to enhance our early stopping algorithm for multi-modal datasets and various GAN variants. We also aim to extend its applicability to other generative models. For additional discussions, please refer to appendix J.



## References

- Dongsheng An, Yang Guo, Na Lei, Zhongxuan Luo, Shing-Tung Yau, and Xianfeng Gu. AE-OT: A new generative model based on extended semi-discrete optimal transport. In *International Conference on Learning Representations*, 2020.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pp. 214–223. PMLR, 2017.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Evan Becker, Parthe Pandit, Sundeep Rangan, and Alyson Fletcher. Instability and local minima in GAN training with kernel discriminators. In *Advances in Neural Information Processing Systems*, volume 35, 2022.
- Giulio Biroli, Tony Bonnaire, Valentin de Bortoli, and Marc Mézard. Dynamical regimes of diffusion models. *arXiv preprint arXiv:2402.18491*, 2024.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019.
- Ricky TQ Chen, Jens Behrmann, David K Duvenaud, and Jörn-Henrik Jacobsen. Residual flows for invertible generative modeling. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Xi Chen, Diederik P. Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel. Variational lossy autoencoder. In *International Conference on Learning Representations*, 2017.
- Rewon Child. Very deep VAEs generalize autoregressive models and can outperform them on images. In *International Conference on Learning Representations*, 2021.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *International Conference on Learning Representations*, 2017.
- Rick Durrett. *Probability: Theory and Examples*, volume 49. Cambridge University Press, 2019.
- Jean-Yves Franceschi, Mike Gartrell, Ludovic Dos Santos, Thibaut Issenhuth, Emmanuel de Bezenac, Mickael Chen, and Alain Rakotomamonjy. Unifying GANs and score-based diffusion as generative particle models. In *Advances in Neural Information Processing Systems*, volume 36, 2023.
- Ruiqi Gao, Yang Song, Ben Poole, Ying Nian Wu, and Diederik P Kingma. Learning energy-based models by diffusion recovery likelihood. In *International Conference on Learning Representations*, 2021.
- Yuan Gao, Yuling Jiao, Yang Wang, Yao Wang, Can Yang, and Shunkang Zhang. Deep generative learning via variational gradient flow. In *International Conference on Machine Learning*, pp. 2093–2101. PMLR, 2019.
- Arnab Ghosh, Viveka Kulharia, Vinay P Nambodiri, Philip HS Torr, and Puneet K Dokania. Multi-agent diverse generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8513–8521, 2018.
- Ian Goodfellow. NIPS 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*, 2017.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, 2014.

- Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, and David Duvenaud. Scalable reversible generative models with free-form continuous dynamics. In *International Conference on Learning Representations*, 2019.
- Xianfeng Gu, Na Lei, and Shing-Tung Yau. Optimal transport for generative models. In *Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging: Mathematical Imaging and Vision*, pp. 1–48. Springer, 2021.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems*, 30, 2017.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Yu-Jui Huang and Yuchong Zhang. GANs as gradient flows that converge. *Journal of Machine Learning Research*, 24(217):1–40, 2023.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pp. 448–456. PMLR, 2015.
- Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35, 2022.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.
- Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible  $1 \times 1$  convolutions. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.
- Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. In *International Conference on Machine Learning*, pp. 1558–1566. PMLR, 2016.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Na Lei, Yang Guo, Dongsheng An, Xin Qi, Zhongxuan Luo, Shing-Tung Yau, and Xianfeng Gu. Mode collapse and regularity of optimal transportation maps. *arXiv preprint arXiv:1902.02934*, 2019.
- Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. MMD GAN: Towards deeper understanding of moment matching network. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Marvin Li and Sitan Chen. Critical windows: Non-asymptotic theory for feature emergence in diffusion models. *arXiv preprint arXiv:2403.01633*, 2024.
- Zinan Lin, Ashish Kheta, Giulia Fanti, and Sewoong Oh. Pacgan: The power of two samples in generative adversarial networks. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

- Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2794–2802, 2017.
- Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018.
- Tu Nguyen, Trung Le, Hung Vu, and Dinh Phung. Dual discriminator generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Albert No, TaeHo Yoon, Kwon Sehyun, and Ernest K Ryu. WGAN with an infinitely wide generator has no spurious stationary points. In *International Conference on Machine Learning*, pp. 8205–8215. PMLR, 2021.
- Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-GAN: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, volume 29, 2016.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *International Conference on Learning Representations*, 2016.
- Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with VQ-VAE-2. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkhäuser, NY*, 55(58–63):94, 2015.
- Antonio Sclocchi, Alessandro Favero, and Matthieu Wyart. A phase transition in diffusion models reveals the hierarchical nature of data. *arXiv preprint arXiv:2402.16991*, 2024.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- Ruoyu Sun, Tiantian Fang, and Alexander Schwing. Towards a better global loss landscape of GANs. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Ugo Tanielian, Thibaut Issenhuth, Elvis Dohmatob, and Jeremie Mary. Learning disconnected manifolds: A no GAN’s land. In *International Conference on Machine Learning*, pp. 9418–9427. PMLR, 2020.
- Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with PixelCNN decoders. In *Advances in Neural Information Processing Systems*, volume 29, 2016.
- Aäron Van Den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning*, pp. 1747–1756. PMLR, 2016.
- Jianwen Xie, Yang Lu, Ruiqi Gao, Song-Chun Zhu, and Ying Nian Wu. Cooperative training of descriptor and generator networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(1):27–45, 2018.
- Yunfei Yang, Zhen Li, and Yang Wang. On the capacity of deep generative networks for approximating distributions. *Neural Networks*, 145:144–154, 2022.

## A Additional Literature Review

In this section, we provide more detailed literature review.

**Generative models.** Learning the generative model based on large amounts of data is a fundamental task in machine learning and statistics. Popular techniques include Variational Autoencoders (Kingma & Welling, 2014; Chen et al., 2017; Razavi et al., 2019; Child, 2021), Generative Adversarial Networks (Goodfellow et al., 2014; Radford et al., 2016; Arjovsky et al., 2017; Gulrajani et al., 2017; Nguyen et al., 2017; Ghosh et al., 2018; Lin et al., 2018; Brock et al., 2019; Karras et al., 2020), flow-based generative models (Dinh et al., 2017; Kingma & Dhariwal, 2018; Chen et al., 2019; Grathwohl et al., 2019), autoregressive models (Van den Oord et al., 2016; Van Den Oord et al., 2016), energy-based models (Xie et al., 2018; Gao et al., 2021) and diffusion models (Ho et al., 2020; Song et al., 2021; Karras et al., 2022). Among these models, GANs’ ability for rapid sampling, unsupervised feature learning and broad applicability makes them the primary focus of this study.

**Practical considerations of GANs.** In the realm of GAN, mode collapse (Goodfellow, 2017) is one of the major challenges which has received a lot of attention. It refers to the situation where the generator produces samples on only a few modes instead of the entire data distribution. The issue of mode collapse has been addressed mainly from three perspectives: modifying the network architecture, designing new objective functions and using normalization techniques. Regarding the network architecture, existing approaches involve increasing the number of generator (Ghosh et al., 2018) or discriminator (Nguyen et al., 2017), using joint architectures (Larsen et al., 2016). From the objective function side, various metrics such as the Wasserstein distance (Arjovsky et al., 2017),  $f$ -divergence (Nowozin et al., 2016), least squares distance (Mao et al., 2017), maximum mean discrepancy (Li et al., 2017) are employed. Normalization techniques such as batch normalization (Ioffe & Szegedy, 2015), layer normalization (Ba et al., 2016) and spectral normalization (Miyato et al., 2018) have also achieved superb empirical performance. Mode mixture (Lei et al., 2019) is another troublesome phenomenon in which the generated samples fall outside the real distribution and are thus unrealistic. Existing approaches include picking generated samples using a rejection sampling method (Tanielian et al., 2020), or generating samples with discontinuous optimal transport rather than deep neural networks (Lei et al., 2019; An et al., 2020; Gu et al., 2021).

**Theoretical frameworks of GANs.** Another line of research approaches mode collapse and mode mixture by establishing theoretical frameworks for better analyzing and optimizing GAN training. These researches fall into two categories: landscape analysis and dynamic analysis. Landscape analysis is static because it simply examines the results of GAN training; it ignores the interaction between the discriminator and generator during training. For instance, Sun et al. (2020) analyzed the landscape of a family of GANs called separable-GAN. They proved that the landscape of separable-GAN has exponentially many bad basins, all of which are deemed as mode-collapse. In contrast, No et al. (2021) demonstrated that Wasserstein GAN with an infinitely broad generator has no spurious stationary points by modeling both the generator and the discriminator using random feature theory. Lei et al. (2019) used results from optimal transport theory to account for mode mixture. Dynamic analysis, on the other hand, considers how the discriminator and generator interact. Becker et al. (2022) suggested the “Isolated Points Model” to explain the causes of GANs’ instability. Another dynamical way of modeling GANs is to regard it as a particle model (Huang & Zhang, 2023; Franceschi et al., 2023). This kind of modeling is used in conjunction with Fokker–Planck equation theories by Huang and Zhang to demonstrate the convergence of GANs to the global stationary point (Huang & Zhang, 2023).

**Relation of GANs and particle models.** There has been an emerging trend in recent years to conceptualize GANs as particle models. We present the interpretation of NSGAN as a particle model in algorithm 1. Huang & Zhang (2023) examined a similar interpretation of vanilla GAN, but did not specifically discuss NSGAN. Gao et al. (2019) used a variational gradient flow approach to analyse GANs, without placing much emphasis on the connection to particle models. Franceschi et al. (2023) unified GANs within the context of particle models and interpret GANs as “interactive particle models”.

**Phased process in diffusion models.** Recently, analogous phase transition phenomena, akin to those elucidated in our paper, have been uncovered in diffusion models. For example, Biroli et al. (2024) showed

that the generative process in diffusion models undergoes a “speciation” transition, revealing data structure from noise, followed by a “collapse” transition, converging dynamics to memorized data points, akin to condensation in a glass phase. Sclocchi et al. (2024) found that the backward diffusion process acting after a time  $t$  is governed by a phase transition at some threshold time, where the probability of reconstructing high-level features suddenly drops and the reconstruction of low-level features evolves smoothly across the whole diffusion process. Li & Chen (2024) studied properties of critical windows that are narrow time intervals in sampling during which particular features of the final image emerge.

## B Choice of Latent Dimension

In this section, we provide the rationale behind our choice of the latent dimension in assumption 2.1. At the population level, Gao et al. (2019) demonstrated that NSGAN minimizes the  $f$ -divergence  $D_f(p_z \| p_{\text{data}})$  with

$$f(u) = (u + 1) \log(u + 1) - 2 \log 2.$$

Let  $\mu$  and  $\nu$  be mutually singular measures on  $\mathbb{R}^n$ , Yang et al. (2022) proved that

$$D_f(\mu \| \nu) = f(0) + f^*(0) > 0,$$

where  $f^*$  stands for the Fenchel conjugate of  $f$ . If the latent dimension is less than  $n$ , then  $g_{\theta\#}p_z$  is supported on a low-dimensional manifold, so that  $g_{\theta\#}p_z$  and  $\nu$  will be mutually singular. Thus there is always a positive gap in  $f$ -divergence between  $g_{\theta\#}p_z$  and  $\nu$ . In other words,  $g_{\theta\#}p_z$  cannot approximate  $\nu$  well even if the GAN model has been trained perfectly. To prevent such inherent misalignment, we assume that the latent dimension always equals  $n$ . Combined with the continuous data augmentation of real-world datasets, we assume that the noise prior  $p_z(\mathbf{z})$  is an  $n$ -dimensional standard Gaussian distribution, denoted as  $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ , where  $n$  is the dimension of real samples.

## C Proofs to Theorems

Here, we aggregate all the theorems presented in the paper and furnish proofs for some of them.

### C.1 Equivalence of NSGAN with Its Particle Model Interpretation

**Theorem C.1.** *The update of  $g_\theta$  via descending the gradient*

$$\nabla_\theta \frac{1}{m} \sum_{i=1}^m \|g_\theta(\mathbf{z}_i) - \hat{\mathbf{Z}}_i\|_2^2$$

*in algorithm 1 is equivalent to descending the gradient*

$$-\nabla_\theta \frac{1}{m} \sum_{i=1}^m \log(d_\omega(g_\theta(\mathbf{z}_i)))$$

*in the original formulation of NSGAN.*

*Proof.* We prove by directly computing the gradient using the chain rule. In fact, we have

$$\begin{aligned} \nabla_\theta \frac{1}{m} \sum_{i=1}^m \|g_\theta(\mathbf{z}_i) - \hat{\mathbf{Z}}_i\|_2^2 &= \frac{2}{m} \sum_{i=1}^m \nabla_\theta g_\theta(\mathbf{z}_i)^\top \cdot (g_\theta(\mathbf{z}_i) - \hat{\mathbf{Z}}_i) \\ &= -\frac{s}{m} \sum_{i=1}^m \nabla_\theta g_\theta(\mathbf{z}_i)^\top \cdot \frac{\nabla d_\omega(\mathbf{Z}_i)}{d_\omega(\mathbf{Z}_i)} \\ &= -s \nabla_\theta \frac{1}{m} \sum_{i=1}^m \log(d_\omega(g_\theta(\mathbf{z}_i))). \end{aligned}$$

This completes the proof. □

## C.2 Properties of Particle Update Dynamics

**Theorem C.2.** Assume that the discriminator is optimal, i.e.,  $d_\omega^*(\mathbf{x}) = p_{\text{data}}(\mathbf{x})/(p_{\text{data}}(\mathbf{x}) + p_g(\mathbf{x}))$ . Denote  $r(\mathbf{x}) = p_{\text{data}}(\mathbf{x})/p_g(\mathbf{x})$ . At a point  $\mathbf{x}$  where  $r(\mathbf{x}) \approx 0$ ,  $\mathbf{x}$  is updated following approximately  $\nabla \log(r(\mathbf{x}))$ . Conversely, when  $r(\mathbf{x}) \gg 1$ ,  $\mathbf{x}$  is updated following approximately  $\nabla(-1/r(\mathbf{x}))$ .

*Proof.* In the following, we abbreviate  $d_\omega^*(\mathbf{x})$  as  $d(\mathbf{x})$ . We then rewrite  $\nabla d(\mathbf{x})/d(\mathbf{x})$  in terms of  $r(\mathbf{x})$ :

$$\begin{aligned} \frac{\nabla d(\mathbf{x})}{d(\mathbf{x})} &= \frac{p_{\text{data}}(\mathbf{x})\nabla p_g(\mathbf{x}) - p_g(\mathbf{x})\nabla p_{\text{data}}(\mathbf{x})}{(p_{\text{data}}(\mathbf{x}) + p_g(\mathbf{x}))p_{\text{data}}(\mathbf{x})} \\ &= \nabla \left( \frac{p_{\text{data}}(\mathbf{x})}{p_g(\mathbf{x})} \right) \cdot \frac{p_g(\mathbf{x})}{p_{\text{data}}(\mathbf{x})(p_{\text{data}}(\mathbf{x}) + p_g(\mathbf{x}))} \\ &= \nabla r(\mathbf{x}) \cdot \frac{1}{r(\mathbf{x})(1 + r(\mathbf{x}))}. \end{aligned}$$

When  $r(\mathbf{x}) \approx 0$ , we have

$$\frac{1}{r(\mathbf{x})(1 + r(\mathbf{x}))} \approx \frac{1}{r(\mathbf{x})}.$$

As a result,

$$\frac{\nabla d(\mathbf{x})}{d(\mathbf{x})} \approx \nabla \log r(\mathbf{x}).$$

When  $r(\mathbf{x}) \gg 1$ , we have

$$\frac{1}{r(\mathbf{x})(1 + r(\mathbf{x}))} \approx \frac{1}{r(\mathbf{x})^2}.$$

Consequently,

$$\frac{\nabla d(\mathbf{x})}{d(\mathbf{x})} \approx \nabla \left( -\frac{1}{r(\mathbf{x})} \right). \quad \square$$

## C.3 Characterization of Measuring-Preserving Maps

**Lemma C.1.** (Durrett, 2019) Let  $\mathbf{X}$  be a random variable taking values on  $\mathbb{R}$  and let  $F_{\mathbf{X}}(x)$  be its CDF. Then

$$F_{\mathbf{X}}^{-1}(\mathcal{U}(0, 1)) \sim \mathbf{X}$$

and

$$F_{\mathbf{X}}(\mathbf{X}) \sim \mathcal{U}(0, 1),$$

where  $\mathcal{U}(0, 1)$  denotes the uniform distribution on  $(0, 1)$ .

**Theorem C.3.** Let  $\Phi(x)$  denotes the cumulative distribution function (CDF) of  $\mathcal{N}(0, 1)$  and let  $\Psi(x)$  be that of  $p_{\text{data}}(x)$ . If  $g$  satisfies  $g_{\#}p_z = p_{\text{data}}$ , then  $g = \Psi^{-1} \circ h \circ \Phi$ , where  $h$  is a measure-preserving map of  $\mathcal{U}(0, 1)$ , i.e., the uniform distribution on  $(0, 1)$ .

*Proof.* We only need to show that  $\Psi \circ g \circ \Phi^{-1}$  is a measure-preserving map of  $\mathcal{U}(0, 1)$ . In fact, by lemma C.1, we have

$$(\Psi \circ g \circ \Phi^{-1})_{\#}\mathcal{U}(0, 1) = (\Psi \circ g)_{\#}p_z = \Psi_{\#}p_{\text{data}} = \mathcal{U}(0, 1). \quad \square$$

## C.4 Steepness of Measure-Preserving Map in 1-Dimension

**Theorem C.4.** Assume that

$$p_{\text{data}} \sim 0.5\mathcal{N}(-\mu, 1) + 0.5\mathcal{N}(\mu, 1).$$

Let  $\Phi(x)$  denote the cumulative distribution function (CDF) of  $\mathcal{N}(0, 1)$  and let  $\Psi(x)$  be that of  $p_{\text{data}}(x)$ . Then  $g(x) := \Psi^{-1}(\Phi(x))$  satisfies

$$g'(0) = \exp(\mu^2/2).$$

*Proof.* By the formula for the derivative of inverse functions, we have

$$g'(0) = \frac{\Phi'(0)}{\Psi'(\Psi^{-1}(\Phi(0)))} = \frac{\Phi'(0)}{\Psi'(0)} = \exp(\mu^2/2). \quad \square$$

### C.5 Steepness of Measure-Preserving Maps in Higher Dimensions

The standard result in (Durrett, 2019) specifically addresses the case of lemma C.2 where  $K = 1$ . And it can be straightforwardly extended to encompass any  $K$ .

**Lemma C.2.** (Durrett, 2019) *Let  $\mathbf{X} \sim \rho(\mathbf{x})d\mathbf{x}$  be a  $n$ -dimensional random vector. Let  $\mathcal{D} \subset \mathbb{R}^n$  satisfy  $\mathbb{P}(\mathbf{X} \in \mathcal{D}) = 1$ . Assume that the map*

$$\varphi: \mathcal{D} = \bigsqcup_{k=1}^K \mathcal{D}_k \rightarrow \mathbb{R}^n$$

*satisfies the following requirements: for each  $1 \leq k \leq K$ ,  $\varphi|_{\mathcal{D}_k}$  is injective and its inverse function is continuously differentiable. Then the probability density function of  $\mathbf{Y} = \varphi(\mathbf{X})$  is*

$$\rho_{\mathbf{Y}}(\mathbf{y}) = \sum_{k=1}^K \rho_{\mathbf{X}}(\varphi^{-1}(\mathbf{y})) \cdot \left| \det(J_{\varphi_k^{-1}}(\mathbf{y})) \right| \cdot \mathbf{1}_{\varphi(\mathcal{D}_k)}(\mathbf{y}).$$

*Equivalently, for any  $\mathbf{x} \in \mathcal{D}$ ,*

$$\rho_{\mathbf{Y}}(\varphi(\mathbf{x})) = \sum_{k=1}^K \rho_{\mathbf{X}}(\mathbf{x}) \cdot |\det(J_{\varphi}(\mathbf{x}))|^{-1} \cdot \mathbf{1}_{\varphi(\mathcal{D}_k)}(\varphi(\mathbf{x})).$$

**Theorem C.5.** *Let  $\nu$  be the truncated Gaussian distribution  $\mathcal{N}_r(\mathbf{0}, \mathbf{I}_n)$  in the  $n$ -dimensional ball  $\mathcal{B}_r(\mathbf{0})$  and assume that  $\hat{\mu}$  is a probability measure with probability density function*

$$\hat{f}(\mathbf{x}) = \frac{1}{Nh} \sum_{i=1}^N \left( \sqrt{1/\pi} \right)^n \exp \left( -\frac{\|\mathbf{x} - \mathbf{x}_i\|_2^2}{h^2} \right).$$

*Suppose that  $g: \mathcal{B}_r(\mathbf{0}) \rightarrow \mathbb{R}^n$  is continuously differentiable and piecewise injective. Then  $\mathcal{S}(g) > M$ , where*

$$M = \delta \cdot h^{1/n} \cdot \sqrt{\frac{\pi}{n}} \cdot \max_{1 \leq i \leq N} \exp \left( \frac{\|\bar{\mathbf{x}} - \mathbf{x}_i\|_2^2}{nh^2} \right).$$

*Here,  $\bar{\mathbf{x}} = \sum_{i=1}^N \mathbf{x}_i/n$ , and  $\delta = \exp(-r^2/2)/\sqrt{2\pi}$ .*

*Proof.* Let  $\mathcal{D}_k$  ( $1 \leq k \leq K$ ) be a partition of  $\mathcal{B}_r(\mathbf{0})$  such that for each  $1 \leq k \leq K$ ,  $g|_{\mathcal{D}_k}$  is injective. We regard  $g$  as the composition of two functions  $g := g_2 \circ g_1$ . Here,  $g_1: \mathcal{B}_r(\mathbf{0}) \rightarrow (0, 1)^n$  satisfies

$$g_1(\mathbf{x}) = g_1(x_1, x_2, \dots, x_n) = (\Phi_r(x_1), \Phi_r(x_2), \dots, \Phi_r(x_n)),$$

where  $\Phi_r(\cdot)$  is the cumulative density function of the 1-dimensional standard Gaussian distribution truncated in  $(-r, r)$ . It is straightforward to show that the derivative of  $\Phi_r$  has a positive lower bound, say,

$$\delta := \frac{1}{\sqrt{2\pi}} \exp \left( -\frac{r^2}{2} \right).$$

Thus  $|\det J_{g_1}(\mathbf{x})| \geq \delta^n$  for any  $\mathbf{x} \in \mathcal{B}_r(\mathbf{0})$ .

By lemma C.1,  $g_{1\#}\nu = \pi$ , where  $\pi$  is the uniform distribution on  $(0, 1)^n$ . In the rest of the proof, we direct our focus to  $g_2: (0, 1)^n \rightarrow \mathbb{R}^n$ , which satisfies  $g_{2\#}\pi = \hat{\mu}$ . Because  $g_2 = g \circ g_1^{-1}$  and  $g$  is injective

on  $\mathcal{D}_i$  ( $1 \leq i \leq N$ ), we conclude that  $g_2$  is injective on  $g_1(\mathcal{D}_k)$  ( $1 \leq k \leq K$ ). By applying lemma C.2 to  $g_2$  and  $g_1(\mathcal{D}_k)$  ( $1 \leq k \leq K$ ), we deduce that for  $\mathbf{y} \in (0, 1)^n$ ,

$$\hat{f}(g_2(\mathbf{y})) = \sum_{k=1}^K \frac{1}{|\det(J_{g_2}(\mathbf{y}))|} \cdot 1_{g_2(g_1(\mathcal{D}_k))}(g_2(\mathbf{y})) \geq \sum_{k=1}^K \frac{1}{|\det(J_{g_2}(\mathbf{y}))|} \cdot 1_{g_1(\mathcal{D}_k)}(\mathbf{y}) = \frac{1}{|\det(J_{g_2}(\mathbf{y}))|}.$$

Let  $\mathcal{B}_R(\mathbf{0})$  be the  $n$ -dimensional open ball centered at the origin with radius  $R = 2 \max_{1 \leq i \leq N} \|\mathbf{x}_i\|_2$ . We consider the point  $\mathbf{y}_0$  satisfying

$$g_2(\mathbf{y}_0) = \arg \max_{\mathbf{x} \in \mathcal{B}_R(\mathbf{0})} \min_{1 \leq i \leq N} \|\mathbf{x} - \mathbf{x}_i\|_2.$$

If there are many of them, we randomly pick one. Let  $\bar{\mathbf{x}} = \sum_{i=1}^N \mathbf{x}_i/n$ . For this  $\mathbf{y}_0$ , we have

$$\hat{f}(g_2(\mathbf{y}_0)) \leq \hat{f}(\bar{\mathbf{x}}) = \frac{1}{Nh} \sum_{i=1}^N \left( \sqrt{1/\pi} \right)^n \exp \left( -\frac{\|\bar{\mathbf{x}} - \mathbf{x}_i\|_2^2}{h^2} \right).$$

Hence

$$|\det(J_{g_2}(\mathbf{y}_0))| \geq \hat{f}(g_2(\mathbf{y}_0))^{-1} \geq h \cdot (\sqrt{\pi})^n \min_{1 \leq i \leq N} \exp \left( \frac{\|\bar{\mathbf{x}} - \mathbf{x}_i\|_2^2}{h^2} \right).$$

Recall that we have  $|\det(J_{g_1}(g_2(\mathbf{y}_0)))| \geq \delta^n$ , where  $\delta = \frac{1}{\sqrt{2\pi}} \exp \left( -\frac{r^2}{2} \right)$ .

Combine the above results and we have

$$|\det(J_g(\mathbf{y}_0))| = |\det(J_{g_2}(\mathbf{y}_0)) \det(J_{g_1}(g_2(\mathbf{y}_0)))| \geq h \cdot (\delta \sqrt{\pi})^n \min_{1 \leq i \leq N} \exp \left( \frac{\|\bar{\mathbf{x}} - \mathbf{x}_i\|_2^2}{h^2} \right).$$

If  $\mathcal{S}(g) < M$ , then by Hadamard's theorem which states that the determinant of a matrix is smaller than or equal to the product of the 2-norms of its column vectors, one can show that

$$|\det(J_g(\mathbf{y}_0))| \leq (M\sqrt{n})^n.$$

Substitute for the expression of  $M$ , we will find that the above two expressions contradict.  $\square$

**Theorem C.6.** Assume that  $p_{data} \sim \mathcal{N}(\mathbf{0}, k_*^2 \mathbf{I}_n)$  and that the discriminator is optimal, i.e., the discriminator consistently provides the precise moving direction for the particle. Then  $k_t$ , the steepness of  $g$  at discrete time step  $t$  satisfies

$$k_{t+1} = k_t + s \left( \frac{1}{k_t^2} - \frac{1}{k_*^2} \right) \cdot \frac{1}{1 + \frac{k_t \varphi(k_t \mathbf{x}_0 / k_*)}{k_* \varphi(\mathbf{x}_0)}},$$

where  $0 \leq t \leq T$ , and  $T$  is the maximum time. Here,  $\varphi$  is the probability density function of  $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ .

*Proof.* Let  $\varphi(\mathbf{x})$  be the probability density function of the  $n$ -dimensional standard Gaussian distribution

$$\varphi(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^n}} \cdot \exp \left( -\frac{1}{2} \mathbf{x}^\top \mathbf{x} \right).$$

Then the probability density function of  $\mathcal{N}(\mathbf{0}, k^2 \mathbf{I}_n)$  is  $\varphi(\mathbf{x}/k)/k$ . Let  $\mathbf{x}_t = k_t \mathbf{x}_0$  denotes the position of the particle at time  $t$ . Here,  $k_t$  represents the steepness of the generator function. We investigate the evolution of the particle subject to the vector field given by  $\nabla d(\mathbf{x})/d(\mathbf{x})$ . Assuming the discriminator is optimal, this process is governed by the following explicit formula (Gao et al., 2019):

$$\mathbf{x}_{t+1} = \mathbf{x}_t - s f''(r(\mathbf{x}_t)) \nabla r(\mathbf{x}_t), \quad t = 1, 2, \dots, T.$$



Here,  $s$  denotes the stepsize,  $T$  is the maximum time,  $f(u) = (u + 1) \log(u + 1) - 2 \log 2$  and

$$r(\mathbf{x}) = \frac{\varphi(\mathbf{x}/k_t)/k_t}{\varphi(\mathbf{x}/k_*)/k_*}$$

is the ratio of the probability density function of  $p_g$  and  $p_{\text{data}}$ . By the formula of  $\varphi(\mathbf{x})$ , we deduce that  $\nabla \varphi(\mathbf{x}) = -\varphi(\mathbf{x})\mathbf{x}$ . Below we compute  $\nabla r(\mathbf{x})$  by the chain rule:

$$\begin{aligned} \nabla r(\mathbf{x}) &= \frac{k_*}{k_t} \cdot \frac{\nabla \varphi(\mathbf{x}/k_t) \cdot \varphi(\mathbf{x}/k_*) - \varphi(\mathbf{x}/k_t) \nabla \varphi(\mathbf{x}/k_*)}{\varphi(\mathbf{x}/k_*)^2} \\ &= \frac{k}{k_t} \left( \frac{1}{k_*^2} - \frac{1}{k_t^2} \right) \frac{\varphi(\mathbf{x}/k_t)}{\varphi(\mathbf{x}/k_*)} \mathbf{x}. \end{aligned}$$

Using  $\mathbf{x}_t = k_t \mathbf{x}_0$ , we derive the following recurrent formula for  $\{k_t\}_{t=0}^T$ :

$$k_{t+1} = k_t + s \left( \frac{1}{k_t^2} - \frac{1}{k_*^2} \right) \cdot \frac{1}{1 + \frac{k_t \varphi(k_t \mathbf{x}_0 / k_*)}{k_* \varphi(\mathbf{x}_0)}}. \quad \square$$

## D Disparity Among Modes Across Different Datasets

### D.1 MNIST

**Proporcessing.** We first transform the images in MNIST by sequentially resizing the images to  $64 \times 64$  pixels, converting them to PyTorch tensors, and normalizing the tensor values to the range of  $[-1, 1]$ .

**Computation.** We calculate the average image tensor for each label based on a set of 10 image tensors sharing the same label. Next, we compute the pairwise distances between these average tensors using the Frobenius norm. The resulting distances are visualized as a heatmap in fig. 10.




Figure 10: Frobenius distances between different modes in MNIST. The tensor of the modes are approximated by taking the average of image tensors that share the same label.

### D.2 Fashion MNIST

**Proporcessing.** We first transform the images in Fashion MNIST by first resizing the images to  $64 \times 64$  pixels, converting them to PyTorch tensors, and normalizing the tensor values to the range of  $[-1, 1]$ .

**Computation.** We calculate the average image tensor for each label based on a set of 10 image tensors sharing the same label. Next, we compute the pairwise distances between these average tensors using the Frobenius norm. The resulting distances are visualized as a heatmap in fig. 11.




Figure 11: Frobenius distances between different modes in Fashion MNIST. The tensor of the modes are approximated by taking the average of image tensors that share the same label.

### D.3 CIFAR-10

**Proporcessing.** We first transform the images in CIFAR-10 by sequentially resizing the images to  $64 \times 64$  pixels, converting them to PyTorch tensors, and normalizing the tensor values to the range of  $[-1, 1]$ .

**Computation.** We calculate the average image tensor for each label based on a set of 10 image tensors sharing the same label. Next, we compute the pairwise distances between these average tensors using the Frobenius norm. The resulting distances are visualized as a heatmap in fig. 12.

## E Detailed Experimental Settings

Our codes are all provided in the supplementary material.

### E.1 Verifying Fitting and Refining

**Methodology.** We demonstrate that the phases of fitting and refining exist in real-world datasets. To do this, we use a classification network  $q(\mathbf{x})$  that takes an image tensor  $\mathbf{x}$  as an input and outputs a 10-dimensional vector,

$$(p_0, p_1, \dots, p_9),$$

where each  $p_i \in [0, 1]$  denotes the likelihood of  $\mathbf{x}$  corresponding to the  $i$ th category (e.g., the 1st category in MNIST corresponds to the handwritten digit 1 and the 2nd category in Fashion MNIST represents pullovers). Our focus gravitates towards those  $p_i$ 's that exhibit significant magnitudes. For discernibility, a threshold  $\tau$  is set to  $10^{-2}$ . In other words, if  $p_i > 10^{-2}$ , then there is a notable probability that  $\mathbf{x}$  belongs to the  $i$ th category. Empirical observations suggest that seldom do more than three  $p_i$ 's surpass the designated threshold. Hence, for any  $\mathbf{x}$ , we may pair  $(i, j)$  when both  $p_i$  and  $p_j$  exceed  $\tau$ . By pairing, the intuition is that such  $\mathbf{x}$  potentially resides *between* modes  $i$  and  $j$ . In scenarios where only a single  $p_i$  surpasses  $\tau$ ,  $i$  is paired with itself, implying that the  $\mathbf{x}$  predominantly belongs to the  $i$ th category. We count the occurrences of the pairings  $(i, j)$  ( $0 \leq i, j \leq 9$ ) in a batch of size 256 and visualize them with heatmaps in fig. 6, fig. 14 and




Figure 12: Frobenius distances between different modes in CIAFR-10. The tensor of the modes are approximated by taking the average of image tensors that share the same label.

fig. 15. In these figures, the value of the entry  $(i, j)$  represents the logarithmically transformed occurrence frequency of pair  $(i, j)$  within a batch, adjusted by one, thereby mitigating the impact of dominant diagonal values on the colorbar.

**Classification networks.** We use the MNIST classification network in MNIST classification network and the Fashion MNIST classification network in Fashion MNIST classification network.

**Number of training runs.** We conducted our experiments at least 50 times and consistently observed similar patterns across all trials. Therefore, we randomly selected two of these experiments to present in this paper.

## E.2 Early Stopping

**Early stopping on 3-dimensional Gaussian mixture.** In this part, our codes borrow heavily from NSGAN. Both the generator and the discriminator are implemented as full-connected neural networks with SGD optimizers. Now we elaborate on how to calculate the threshold defined in algorithm 2. The threshold is given by  $k_s \cdot \sqrt{2}/(2h)$ . We set  $k_s = 2$ , the distance between two modes in the 3-dimensional Gaussian mixture dataset. For  $h$ , it equals  $\sqrt{2}\sigma$ , where  $\sigma$  is the standard variation of every Gaussian component, which is  $\sqrt{0.0125}$  in our setting. Therefore the threshold is

$$2 \times \sqrt{2}/(2 \times \sqrt{2} \times \sqrt{0.0125}) \approx 8.9.$$

As for the warm-up training iteration parameter  $N_w$ , we set it to 50.

**Early stopping on MNIST.** In this part, our generator and discriminator architectures borrow heavily from NSGAN on MNIST. Both the generator and the discriminator are implemented as convolutional neural networks with Adam optimizers. Now we elaborate on how to calculate the threshold defined in algorithm 2. The threshold is given by  $k_s \cdot \sqrt{2}/(2h)$ . We set  $k_s = 33.7$ , the distance between two farthest modes in MNIST (please refer to appendix D). For  $h$ , it equals  $\sqrt{2}\sigma$ , where  $\sigma^2$  is derived as follows. We first compute the population variance of the images from each label, arriving at 10 values. Then we compute their average value, and divide this value by  $64 \times 64 \times 1$ , i.e., the total number of dimensions. Therefore the threshold is

$$33.7 \times \sqrt{2}/(2 \times \sqrt{2} \times \sqrt{0.33/64^2}) \approx 1877.$$

As for the warm-up training iteration parameter  $N_w$ , we set it to 20.

**Early stopping on fashion MNIST.** In this part, our generator and discriminator architectures borrow heavily from NSGAN on Fashion MNIST. Both the generator and the discriminator are implemented as convolutional neural networks with Adam optimizers. Now we elaborate on how to calculate the threshold defined in algorithm 2. The threshold is given by  $k_s \cdot \sqrt{2}/(2h)$ . We set  $k_s = 48.1$ , the distance between two farthest modes in Fashion MNIST (please refer to appendix D). For  $h$ , it equals  $\sqrt{2}\sigma$ , where  $\sigma^2$  is derived as follows. We first compute the population variance of the images from each label, arriving at 10 values. Then we compute their average value, and divide this value by  $64 \times 64 \times 1$ , i.e., the total number of dimensions. Therefore the threshold is

$$48.1 \times \sqrt{2}/(2 \times \sqrt{2} \times \sqrt{0.33/64^2}) \approx 2679.$$

As for the warm-up training iteration parameter  $N_w$ , we set it to 50.

**Early stopping on CIFAR-10.** In this part, our generator and discriminator architectures borrow heavily from NSGAN on CIFAR-10. Both the generator and the discriminator are implemented as convolutional neural networks with Adam optimizers. Now we elaborate on how to calculate the threshold defined in algorithm 2. The threshold is given by  $k_s \cdot \sqrt{2}/(2h)$ . We set  $k_s = 38.0$ , the distance between two farthest modes in CIFAR-10 (please refer to appendix D). For  $h$ , it equals  $\sqrt{2}\sigma$ , where  $\sigma^2$  is derived as follows. We first compute the population variance of the images from each label, arriving at 10 values. Then we compute their average value, and divide this value by  $64 \times 64 \times 3$ , i.e., the total number of dimensions. Therefore the threshold is

$$38.0 \times \sqrt{2}/(2 \times \sqrt{2} \times \sqrt{0.23/(64^2 \times 3)}) \approx 4391.$$

As for the warm-up training iteration parameter  $N_w$ , we set it to 50.

**Number of training runs.** On all of the datasets mentioned above, we conducted our experiments at least 100 times. We observed similar patterns across all trials, although the point at which the GANs collapsed varied. Therefore, we choose to present those that collapsed before a certain threshold to ensure consistency in our reported results. It is important to note that the generated samples eventually collapsed in our experiments, either sooner or later, without contradicting the findings in our paper.

## F Additional Experimental Results

### F.1 Overfitting of the Discriminator

In this subsection, we elaborate on the optimal discriminator’s behavior outlined in section 5.1. We consider the following synthetic dataset

$$p_{\text{data}} \sim \frac{1}{4}\mathcal{N}([1, 1], 0.0125\mathbf{I}_2) + \frac{1}{4}\mathcal{N}([1, -1], 0.0125\mathbf{I}_2) + \frac{1}{4}\mathcal{N}([-1, 1], 0.0125\mathbf{I}_2) + \frac{1}{4}\mathcal{N}([-1, -1], 0.0125\mathbf{I}_2),$$

and train the discriminator until optimal. We plot the values of the optimal discriminator in fig. 13. We observe that the discriminator values are close to 0.5 in the central regions of the modes and vanish in the regions far from the modes. Between them, the discriminator values smoothly change from 0.5 to 0.




Figure 13: The values of the optimal discriminator. The discriminator values are close to 0.5 in the central regions of the modes (i.e.,  $[\pm 1, \pm 1]$ ) and vanish in the regions far from the modes. Between them, the discriminator values smoothly change from 0.5 to 0.

## F.2 Verifying Fitting and Refining

**Annotated heatmaps for MNIST.** We verify the existence of fitting and refining on MNIST. Annotated heatmaps are employed to track the evolution of pairings  $(i, j)$  occurrence within batches of size 256. The values depicted in these heatmaps represent the logarithm of occurrence counts plus 1, with darker colors indicating higher values. Each heatmap includes epoch numbers ranging from 0 to 38 displayed at the bottom. Initially, the heatmap has few nonzero entries, indicating limited sample diversity during the fitting phase. As training advances, more entries became nonzero, reflecting a broader distribution of generated samples across the mode space. Notably, the values of off-diagonal entries signifies the severity of mode mixture, which gradually decrease over the course of training, validating the refining phase. However, the issue of mode mixture persists even at the end of refining. By the 36th epoch, the heatmap only has two nonzero entries, suggesting the collapsing phase, where the generated samples become less diverse and concentrate around few modes. These observations provide empirical evidence for our proposed three phases of GAN training.

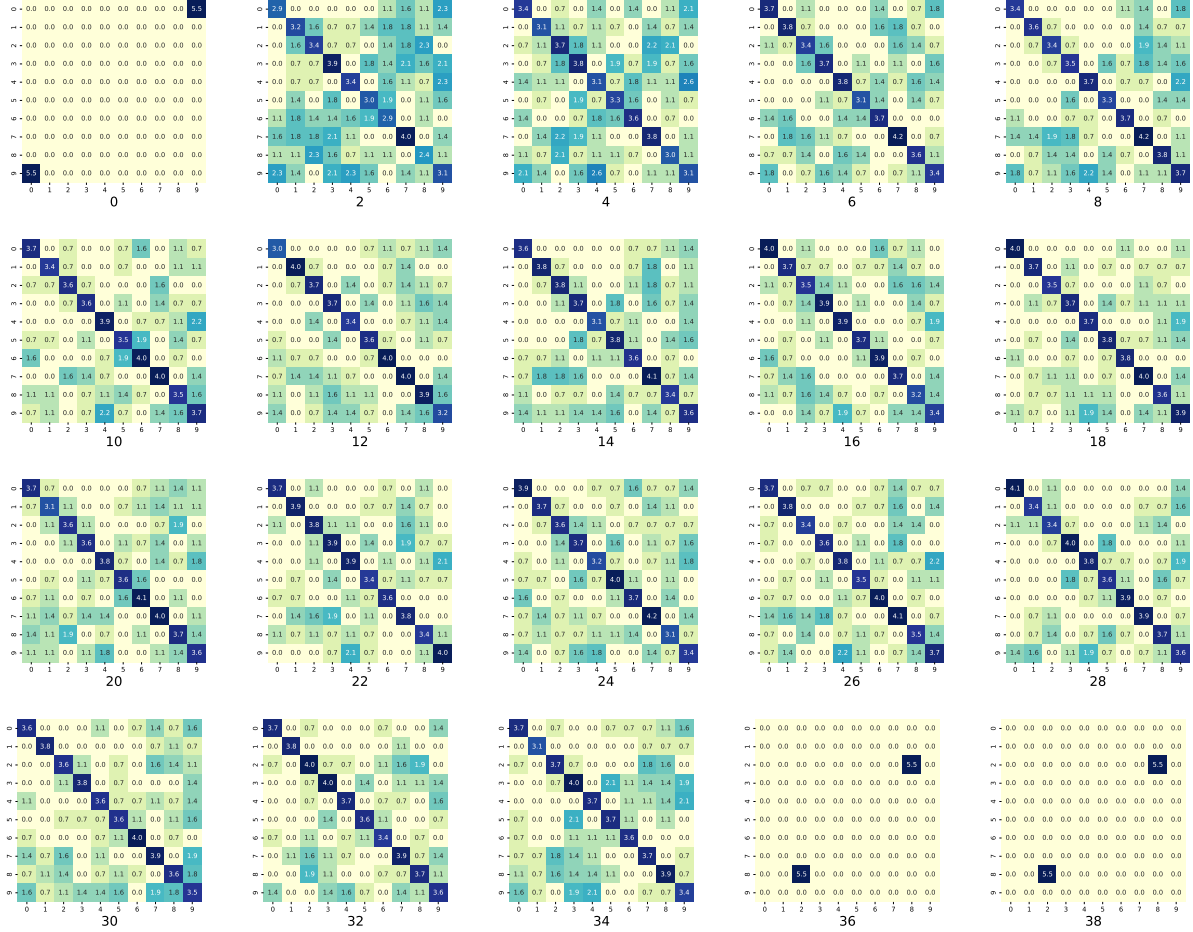


Figure 14: Annotated heatmaps for verifying fitting and refining in MNIST. The values are the logarithm of the occurrence of pairings  $(i, j)$  plus 1 in a batch of size 256. Darker colors indicate higher values. The epochs, ranging from 0 to 38, are displayed at the bottom of each heatmap. Initially, there are few nonzero entries, suggesting limited sample diversity. As training progresses, more entries become nonzero, indicating wider sample distribution across mode space, which corresponds to the fitting phase. Off-diagonal entries reflect mode mixture, which diminishes over training, confirming the refining phase. Remarkably, mode mixture persists even at the closure of the refining phase. Note that by the 36th epoch, only two entries remain nonzero, indicating the collapsing phase.

**Verifying fitting and refining in Fashion MNIST.** We verify the existence of fitting and refining on Fashion MNIST using annotated heatmaps. The heatmap values are the logarithm of pairings  $(i, j)$  occurrence plus 1 in batches of size 256, with darker colors indicating higher values. Each epoch is divided into 5 collections of batches, denoted as  $e, b$  where  $e$  is the epoch and  $b$  is the batch collection within the epoch. Initially, there are only two nonzero entries, which suggests limited sample diversity. As training progresses, more entries become nonzero, indicating a broader sample distribution across the mode space during the fitting phase. Notably, unlike MNIST, the phases of fitting and refining in Fashion MNIST occur quickly, evidenced by the rapid stabilization of off-diagonal values. It is important to note that the large values in some off-diagonal entries do not necessarily imply severe mode mixture. For example, “T-shirt”, “Pullover”, and “Shirt” are frequently confused in Fashion MNIST classification tasks.

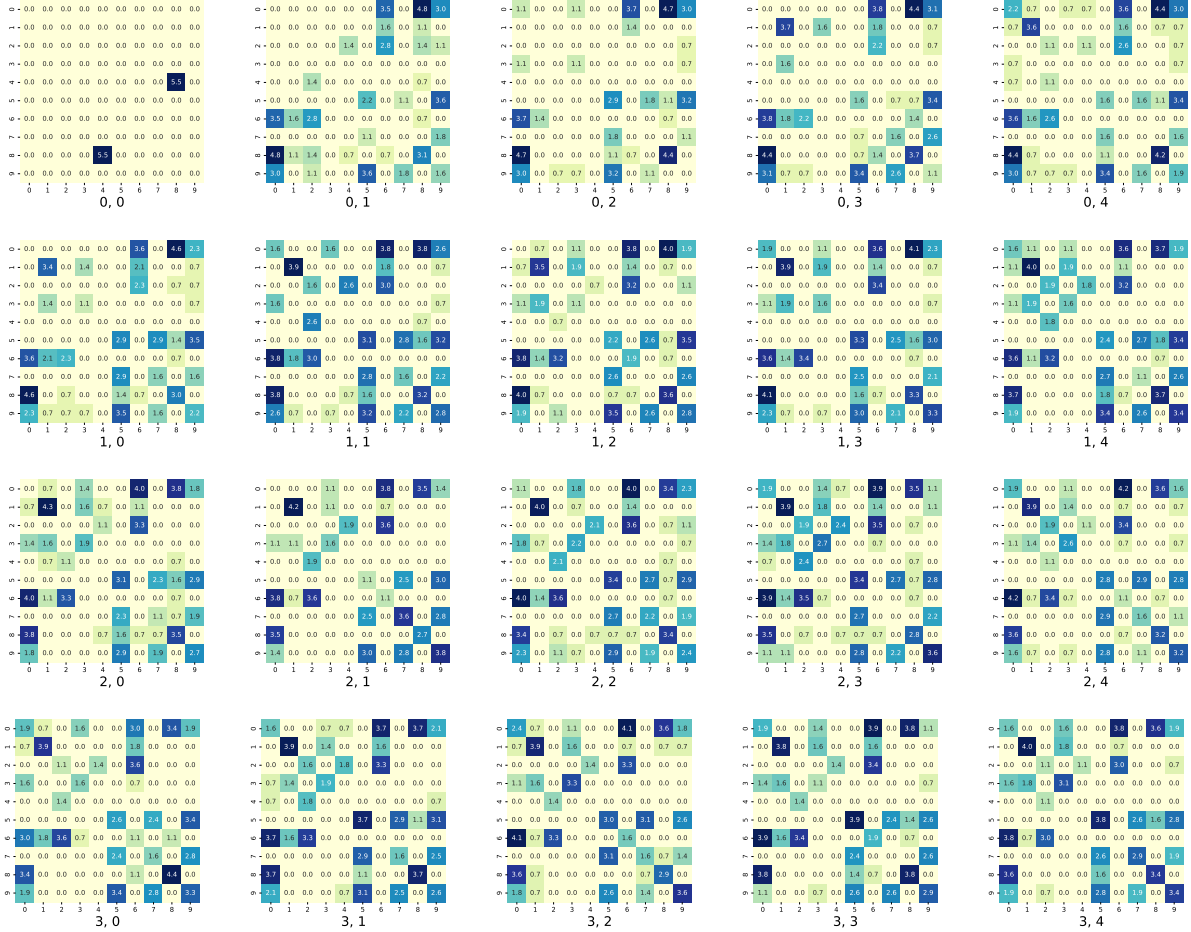


Figure 15: Annotated heatmaps for verifying fitting and refining in Fashion MNIST. The labels 0 to 9 mean “T-shirt/top”, “Trouser”, “Pullover”, “Dress”, “Coat”, “Sandal”, “Shirt”, “Sneaker”, “Bag”, and “Ankle boot”, respectively. The values are the logarithm of the occurrence of pairings  $(i, j)$  plus 1 in a batch of size 256. Darker colors indicate higher values. Each epoch is equally divided into 5 collection of batches. The label “ $e, b$ ” at the bottom of each heatmap denotes the  $b$ th collection within the  $e$ th epoch. Therefore, the heatmaps displayed are for the first 4 epochs only. Initially, the few nonzero entries indicate limited sample diversity. As training progresses, more entries became nonzero, reflecting a broader sample distribution across the mode space, which corresponds to the fitting phase. Unlike MNIST, the phases of fitting and refining in Fashion MNIST take place rapidly because the off-diagonal values stabilize quickly. It is important to note that the large values of some off-diagonal entries do not necessarily imply severe mode mixture; for instance, “T-shirt”, “Pullover”, and “Shirt” are often confused in Fashion MNIST classification tasks.

### F.3 Early Stopping on Gaussian Mixture: More Training Runs

We train NSGAN on the 3-dimensional Gaussian mixture dataset and record down  $\|\nabla d_\omega/d_\omega\|_2$  each epoch until the maximum specified epochs. The threshold is set to 8.9 for reasons in appendix E.2. Instead of halting training when this norm exceeds the threshold, we opt to continue training, which allows us to assess the sample quality both before and after the stop.

In fig. 16, we present the tendency of  $\|\nabla d_\omega/d_\omega\|_2$  against epochs in four trials. We consistently observe a pattern where these values initially increase, then decrease to nearly zero, subsequently rise above the threshold of 8.9, and continue to increase thereafter. This phenomenon can be empirically understood as follows: during the fitting phase, particles are updated to relatively large distances to explore modes; in the refining phase, most particles stabilize near the modes, while the remainder converge to the modes. Finally, in the collapsing phase discussed in section 5.1, particles are pushed away from the modes. The generated samples in the four trials are displayed in fig. 17, fig. 18, fig. 19, and fig. 20, respectively.




Figure 16: The values of  $\|\nabla d_\omega/d_\omega\|_2$  against epochs in four trials. We mark the stopping epochs by red crosses. The patterns of  $\|\nabla d_\omega/d_\omega\|_2$  values across epochs are consistent in all the four training runs. Initially, these values increase, then decrease to nearly zero, and subsequently rise above the threshold of 8.9, continuing to increase thereafter.



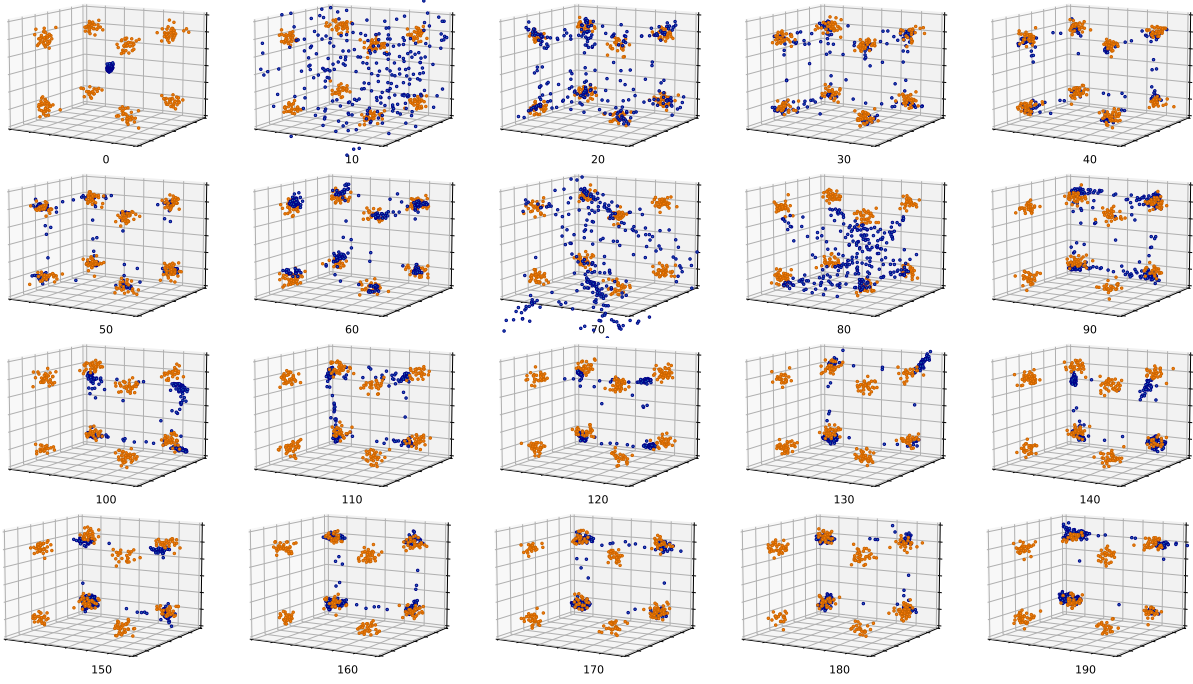


Figure 17: Generated samples in trial 0. The stopping point is the 76th epoch. Before the stopping point, generated samples demonstrate notable quality, particularly evident around the 50th epoch. As training approaches the stopping point, they start to deviate from the modes. After the stopping, they undergo a gradual deterioration, ultimately collapsing to approximately half of the 8 modes.

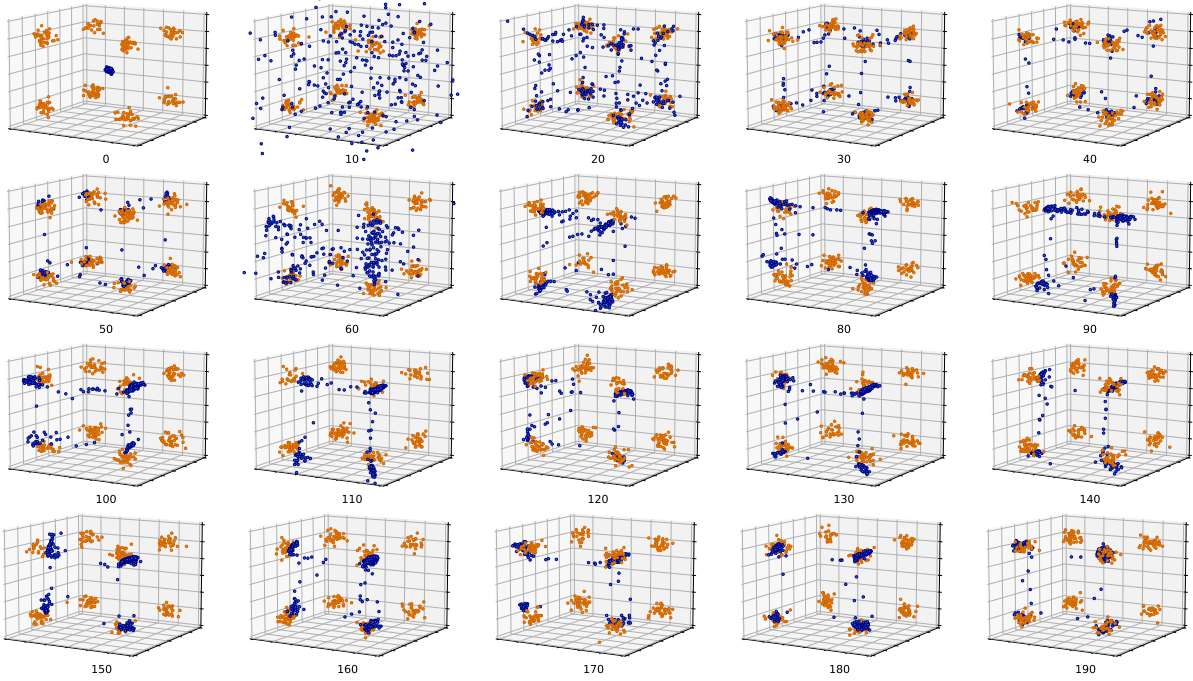


Figure 18: Generated samples in trial 1. The stopping point is the 60th epoch. Before the stopping point, generated samples demonstrate notable quality, particularly evident around the 50th epoch. As training approaches the stopping point, they start to deviate from the modes. After the stopping, they undergo a gradual deterioration, ultimately collapsing to approximately half of the 8 modes.



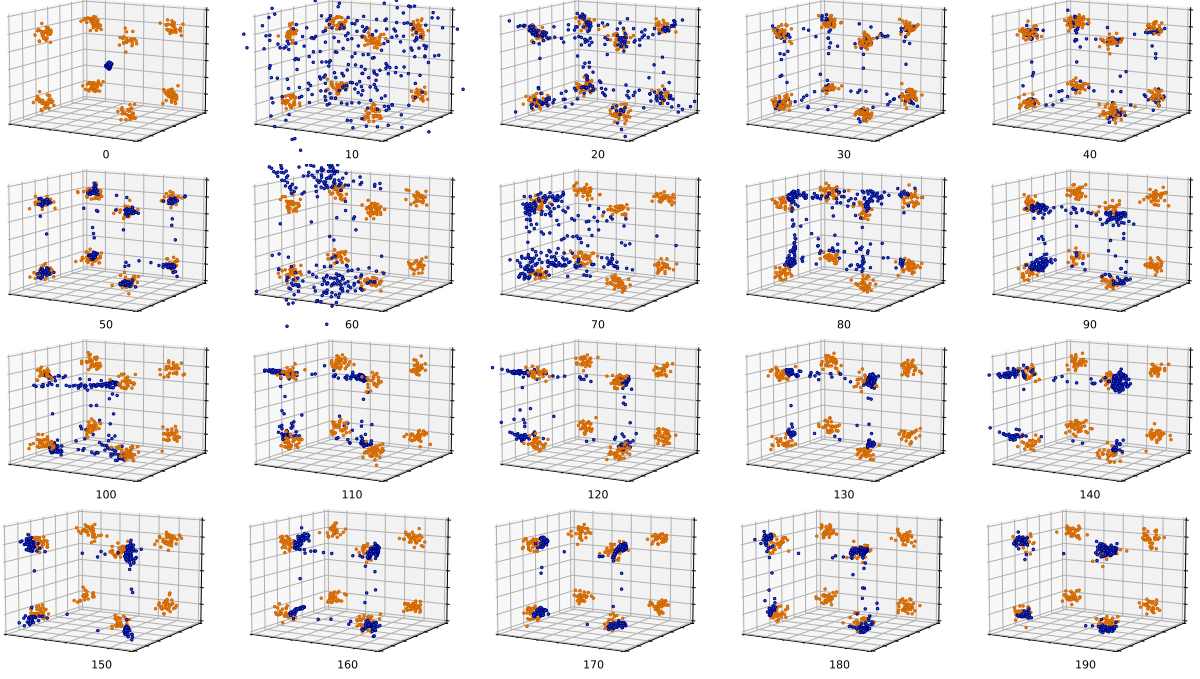


Figure 19: Generated samples in trial 2. The stopping point is the 67th epoch. Before the stopping point, generated samples demonstrate notable quality, particularly evident around the 50th epoch. As training approaches the stopping point, they start to deviate from the modes. After the stopping, they undergo a gradual deterioration, ultimately collapsing to approximately half of the 8 modes.

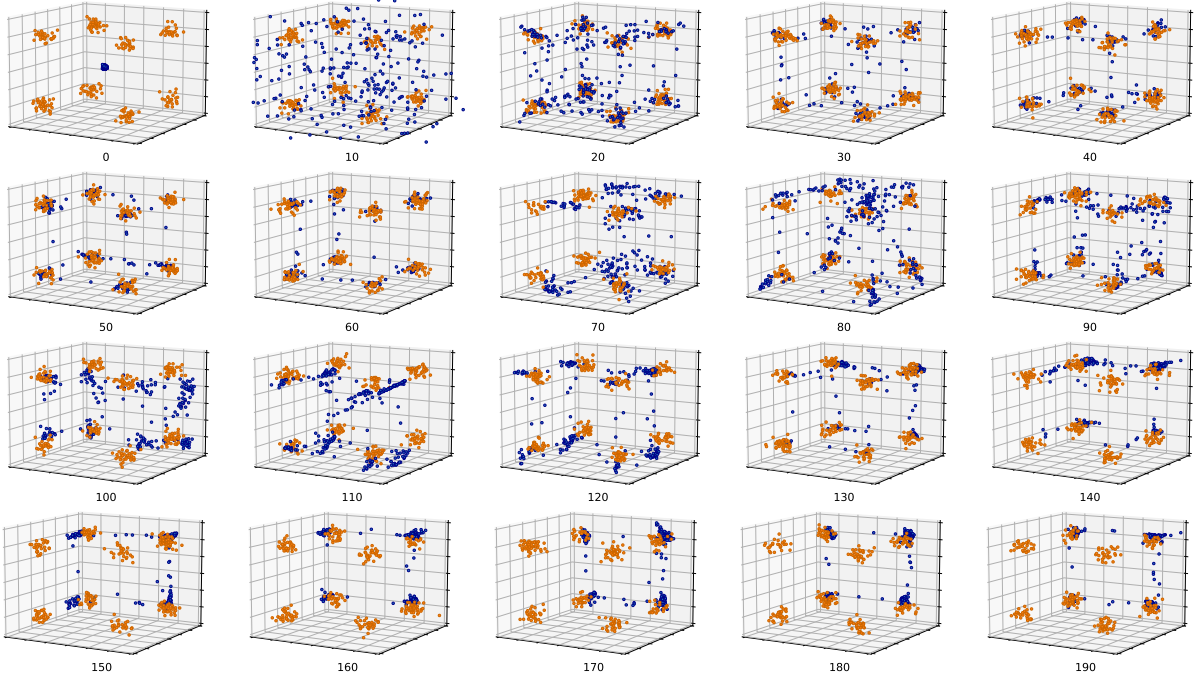


Figure 20: Generated samples in trial 3. The stopping point is the 77th epoch. Before the stopping point, generated samples demonstrate notable quality, particularly evident around the 60th epoch. As training approaches the stopping point, they start to deviate from the modes. After the stopping, they undergo a gradual deterioration, ultimately collapsing to approximately half of the 8 modes.

#### F.4 Early Stopping on MNIST: More Training Runs

We train NSGAN on MNIST and record down  $\|\nabla d_\omega/d_\omega\|_2$  each epoch until the maximum specified epochs. The threshold is set to 1877 for reasons in appendix E.2. Instead of halting training when this norm exceeds the threshold, we opt to continue training, which allows us to assess the sample quality both before and after the stop.

In fig. 21, we present the tendency of  $\|\nabla d_\omega/d_\omega\|_2$  against epochs in four trials. We consistently observe a pattern where these values show a smooth initial increase, followed by turbulent fluctuations, culminating in a sudden surge surpassing the threshold of 1877, and then dropping to approximately zero thereafter. The generated samples in the four trials are displayed in fig. 22, fig. 23, fig. 24, and fig. 25, respectively. Note that these figures should be interpreted from left to right, where each column of images represents the generated samples from the epoch indicated at the bottom of the column.




Figure 21: The values of  $\|\nabla d_\omega/d_\omega\|_2$  against epochs in four trials. We mark the stopping epochs by red crosses. The patterns of  $\|\nabla d_\omega/d_\omega\|_2$  values across epochs are consistent across in four training runs. The values exhibit a smooth initial increase, followed by turbulent fluctuations, culminating in a sudden surge above the 1877 threshold, before dropping to approximately zero.




Figure 22: Generated images in trial 0. The stopping point is the 41st epoch. Before reaching the stopping point, the generated samples exhibit good quality, notably around the 30th epoch. Shortly after the stop, they are contaminated with noise. Then they gradually regain clarity and ultimately collapse to one of the modes.




Figure 23: Generated images in trial 1. The stopping point is the 33rd epoch. Before reaching the stopping point, the generated samples exhibit good quality, notably around the 30th epoch. Shortly after the stop, they are contaminated with noise. Then they gradually regain clarity and ultimately collapse to one of the modes.




Figure 24: Generated images in trial 2. The stopping point is the 57th epoch. Before reaching the stopping point, the generated samples exhibit good quality, notably around the 55th epoch. Shortly after the stop, they are contaminated with noise. Then they gradually regain clarity and ultimately collapse to one of the modes.




Figure 25: Generated images in trial 3. The stopping point is the 36th epoch. Before reaching the stopping point, the generated samples exhibit good quality, notably around the 35th epoch. Shortly after the stop, they are contaminated with noise. Then they gradually regain clarity and ultimately collapse to one of the modes.

### F.5 Early Stopping on Fashion MNIST: More Training Runs

We train NSGAN on Fashion MNIST and record down  $\|\nabla d_\omega/d_\omega\|_2$  each epoch until the maximum specified epochs. The threshold is set to 2679 for reasons in appendix E.2. Instead of halting training when this norm exceeds the threshold, we opt to continue training, which allows us to assess the sample quality both before and after the stop.

In fig. 26, we present the tendency of  $\|\nabla d_\omega/d_\omega\|_2$  against epochs in four trials. We consistently observe a pattern where these values exhibit a smooth initial increase, followed by turbulent fluctuations, culminating in a sudden surge to surpass the threshold of 2679. Note that in the phases of fitting and refining,  $\|\nabla d_\omega/d_\omega\|_2$  may exceed the threshold as well. This situation justifies the necessity of introducing a warm-up training iteration parameter  $N_w$  in algorithm 2. The generated samples in the four trials are displayed in fig. 27, fig. 28, fig. 29, and fig. 30, respectively. Note that these figures should be interpreted from left to right, where each column of images represents the generated samples from the epoch indicated at the bottom of the column.

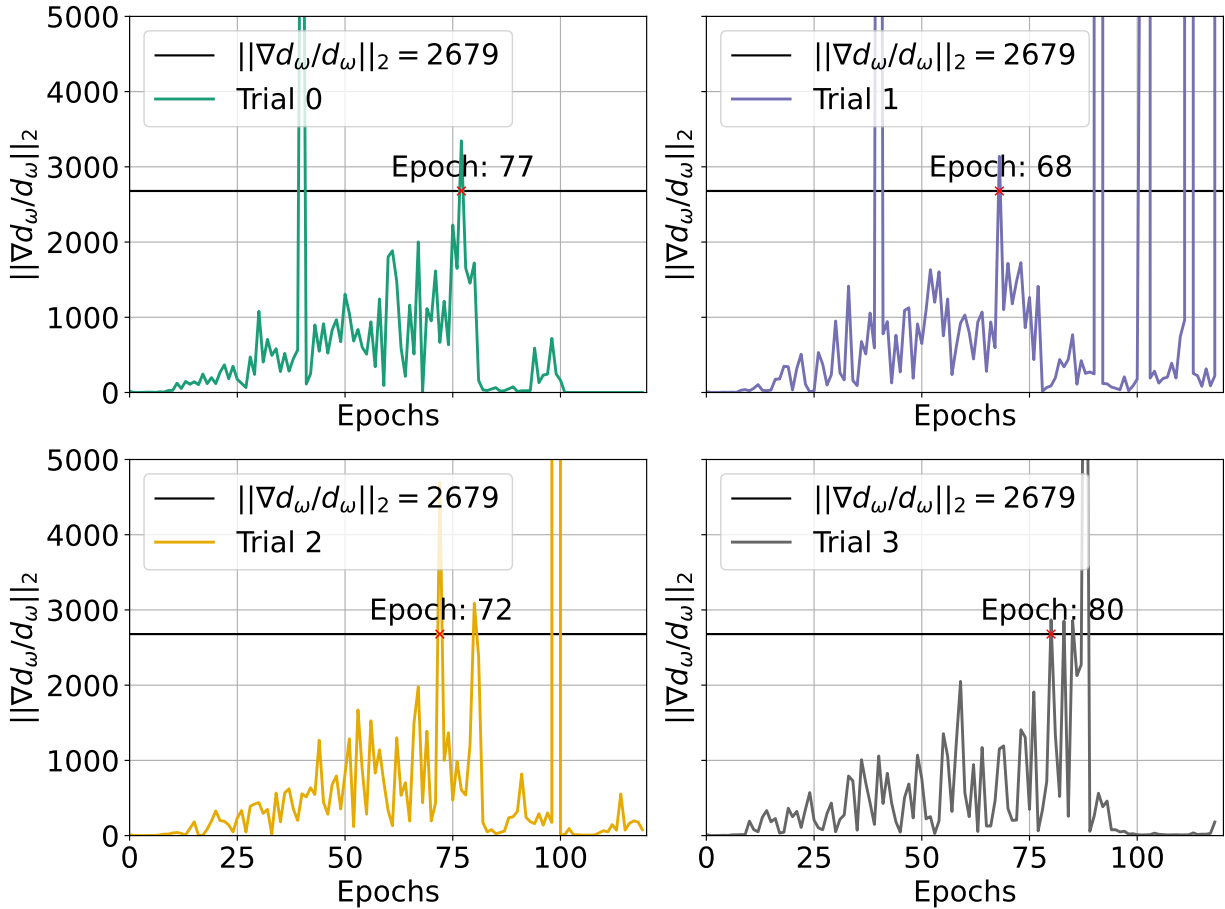


Figure 26: The values of  $\|\nabla d_\omega/d_\omega\|_2$  against epochs in four trials. We mark the stopping epochs by red crosses. Note that only the values of the epochs after  $N_w = 50$  (where  $N_w$  is the warm-up training iteration parameter defined in algorithm 2) that exceed 2679 may be considered as potential stopping points. The patterns of  $\|\nabla d_\omega/d_\omega\|_2$  values across epochs are consistent in all four training runs. Initially, there is a gradual rise, followed by turbulent fluctuations, and eventually a sharp surge above the 2679 threshold.






Figure 27: Generated images in trial 0. The stopping point is the 77th epoch. Before the stopping point, the generated samples demonstrate good quality, particularly evident around the 66th epoch. Shortly after surpassing this point, they become contaminated by noise. After that, there is a gradual recovery of clarity, ultimately resulting in convergence towards two specific modes.



Figure 28: Generated images in trial 1. The stopping point is the 68th epoch. Before the stopping point, the generated samples demonstrate good quality, particularly evident around the 66th epoch. Shortly after surpassing this point, they become contaminated by noise. After that, there is a gradual recovery of clarity, ultimately resulting in convergence towards three specific modes.




Figure 29: Generated images in trial 2. The stopping point is the 72nd epoch. Before the stopping point, the generated samples demonstrate good quality, particularly evident around the 66th epoch. Shortly after surpassing this point, they become contaminated by noise. After that, there is a gradual recovery of clarity, ultimately resulting in convergence towards four specific modes.




Figure 30: Generated images in trial 3. The stopping point is the 80th epoch. Before the stopping point, the generated samples demonstrate good quality, particularly evident around the 78th epoch. Shortly after surpassing this point, they become contaminated by noise. After that, there is a gradual recovery of clarity, ultimately resulting in convergence towards two specific modes.

### F.6 Early Stopping on CIFAR-10: More Training Runs

We train NSGAN on CIFAR-10 and record down  $\|\nabla d_\omega/d_\omega\|_2$  each epoch until the maximum specified epochs. The threshold is set to 4391 for reasons in appendix E.2. Instead of halting training when this norm exceeds the threshold, we opt to continue training, which allows us to assess the sample quality both before and after the stop.

In fig. 31, we present the tendency of  $\|\nabla d_\omega/d_\omega\|_2$  against epochs in four trials. We consistently observe a pattern where these values exhibit a smooth initial increase, followed by turbulent fluctuations, culminating in a sudden surge to surpass the threshold of 4391. The generated samples in the four trials are displayed in fig. 32, fig. 33, fig. 34, and fig. 35, respectively. Note that these figures should be interpreted from left to right, where each column of images represents the generated samples from the epoch indicated at the bottom of the column.




Figure 31: The values of  $\|\nabla d_\omega/d_\omega\|_2$  against epochs in four trials. We mark the stopping epochs by red crosses. The patterns of  $\|\nabla d_\omega/d_\omega\|_2$  values across epochs are consistent in all four training runs. The values show a gradual rise at the beginning, followed by turbulent fluctuations, and then suddenly surge above the threshold of 4391.





Figure 32: Generated images in trial 0. The stopping point is the 76th epoch. Before the stopping point, the generated samples demonstrate good quality, particularly evident around the 72th epoch. Shortly after surpassing this point, they become contaminated by noise. Afterwards, they regain clarity, but transition between different modes.



Figure 33: Generated images in trial 1. The stopping point is the 93rd epoch. Before the stopping point, the generated samples demonstrate good quality, particularly evident around the 78th epoch. Shortly after surpassing this point, they become contaminated by noise. Afterwards, they regain clarity, but transition between different modes.





Figure 34: Generated images in trial 2. The stopping point is the 68th epoch. Before the stopping point, the generated samples demonstrate good quality, particularly evident around the 54th epoch. Approaching the stopping point, they become contaminated by noise. Afterwards, they regain clarity, but transition between different modes.

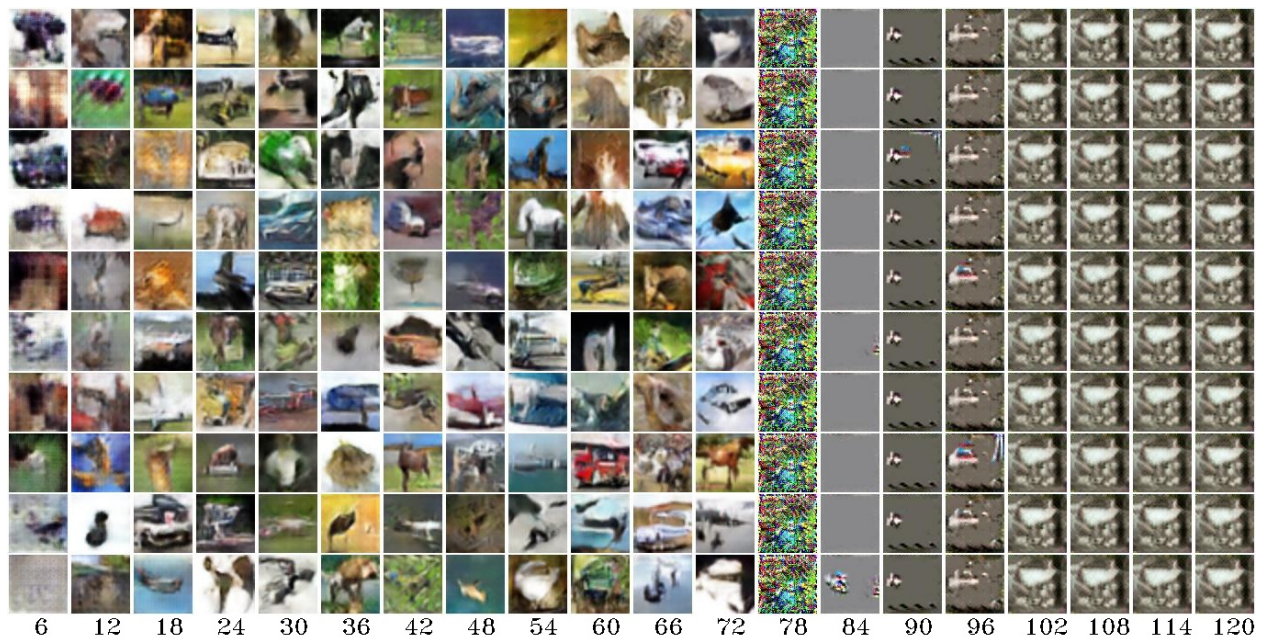


Figure 35: Generated images in trial 3. The stopping point is the 83rd epoch. Before the stopping point, the generated samples demonstrate good quality, particularly evident around the 66th epoch. Shortly after surpassing this point, they become contaminated by noise. Afterwards, they regain clarity and converge towards one specific mode.

### F.7 Relationship between Our Early Stopping Metric and the FID Score

In this subsection, we present the comparison between our early stopping metric  $\|\nabla d_\omega/d_\omega\|_2$  (which we will refer to as the “metric” throughout this subsection) and the FID score. We plot the metric and the FID score in the same figure with shared  $x$ -axis. The results of four trials for MNIST, Fashion MNIST and CIFAR-10 are shown respectively in fig. 36, fig. 37, and fig. 38.

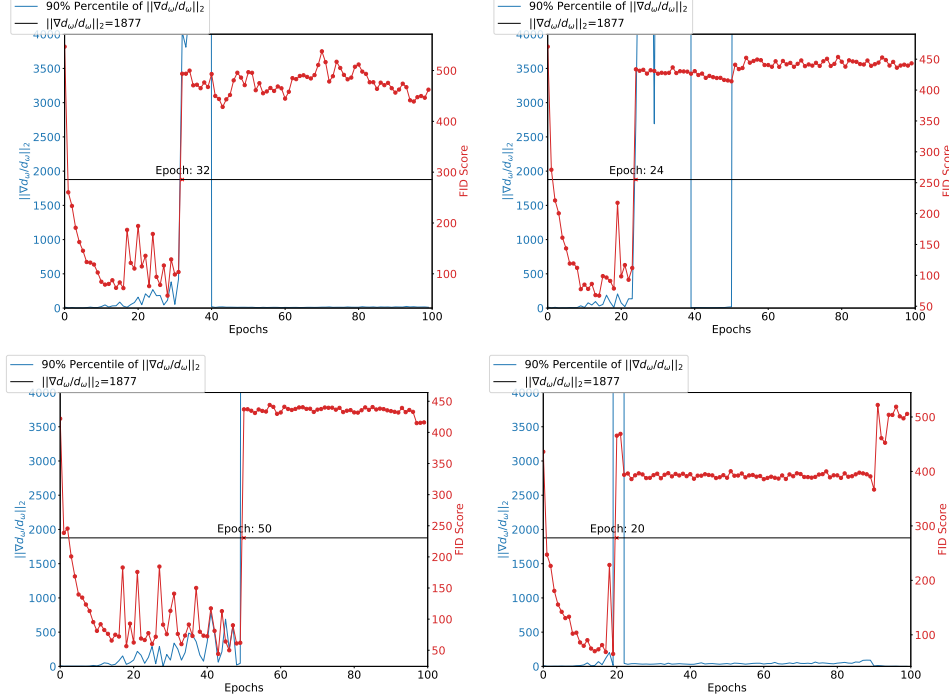



Figure 36: The tendency of  $\|\nabla d_\omega/d_\omega\|_2$  and the FID scores on **MNIST in four trials**. Initially, the metric shows a smooth rise, accompanied by a steady decrease in the FID score, indicating the phase of fitting and refining. Subsequently, the metric exhibits turbulent evolution, and the FID score also rises, signaling the conclusion of refining and the prelude of collapsing. Then both the metric and FID score soar. Consequently, the FID score remains consistently high, indicating the collapsing phase.






Figure 37: The tendency of  $\|\nabla d_\omega/d_\omega\|_2$  and the FID scores on **Fashion MNIST in four trials**. Initially, the metric shows a smooth rise, accompanied by a steady decrease in the FID score, indicating the phase of fitting and refining. Subsequently, the metric exhibits turbulent evolution, but the FID score remains steady. This aligns with our empirical observations in appendix F.5, where generated images demonstrate good quality despite the rapid transition of particles from modes to modes. Then both the metric and FID score soar. Consequently, the FID score remains consistently high, indicating the collapsing phase.

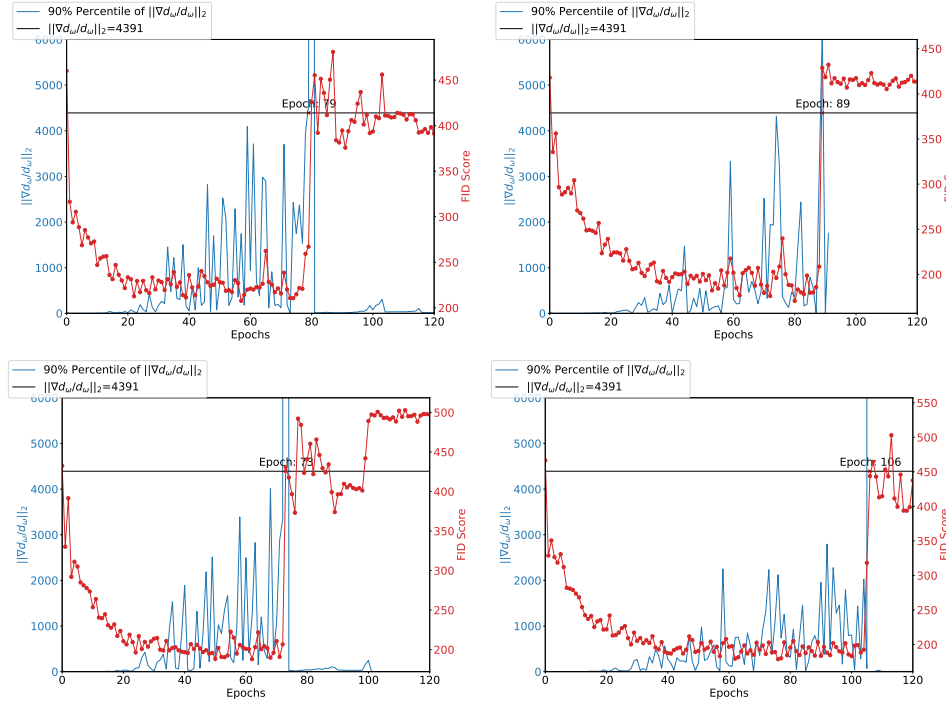


Figure 38: The tendency of  $\|\nabla d_\omega/d_\omega\|_2$  and the FID scores on **CIFAR-10 in four trials**. Initially, the metric shows a smooth rise, accompanied by a steady decrease in the FID score, indicating the phase of fitting and refining. Subsequently, the metric exhibits turbulent evolution, but the FID score remains steady. This aligns with our empirical observations in appendix F.6, where generated images demonstrate good quality despite the rapid transition of particles from modes to modes. Then both the metric and FID score soar. Consequently, the FID score remains consistently high, indicating the collapsing phase.



### F.8 Impact on the Early Stopping Metric after Applying Techniques to Mitigate Mode Collapse

In this subsection, we validate our early stopping metric’s effectiveness by demonstrating that injecting noise into the intermediate layers of the discriminator combats mode collapse and pushes back the metric.

**Experimental setup.** We devise two generator models of identical architecture and implement two discriminators, one adhering to the original design (which we will refer to as “noise-free”) and the other modified to incorporate Gaussian noise with a standard deviation of 0.1 before forwarding the input to the subsequent layer (which we will refer to as “noised”). Both generators and discriminators are initialized using the same random seed. During training, the four networks are concurrently trained, with each generator paired with a discriminator. We present the generated images of the two models on Fashion MNIST in fig. 39 and histograms of  $\|\nabla d_\omega/d_\omega\|_2$  in fig. 40.

**Results.** The noise-free GAN collapses at the 54th epoch, while the noised GAN consistently generates high-quality images. The introduction of noise results in an overall decrease in the  $\|\nabla d_\omega/d_\omega\|_2$  compared to its noise-free counterpart before the 54th epoch. After the 54th epoch, the opposite trend is observed, attributed to the vanishing of  $\|\nabla d_\omega/d_\omega\|_2$  in the noise-free GAN. This indicates that the strategy of injecting noise to mitigate mode collapse leads to an overall decrease in our proposed metric, thereby validating the effectiveness of the metric.





Figure 39: The generated images from the noise-free GAN and the noised GAN. **Upper:** Noise-free GAN. **Lower:** Noised GAN. The noise-free GAN collapses at the 54th epoch, whereas the noised GAN consistently produces high-quality images.



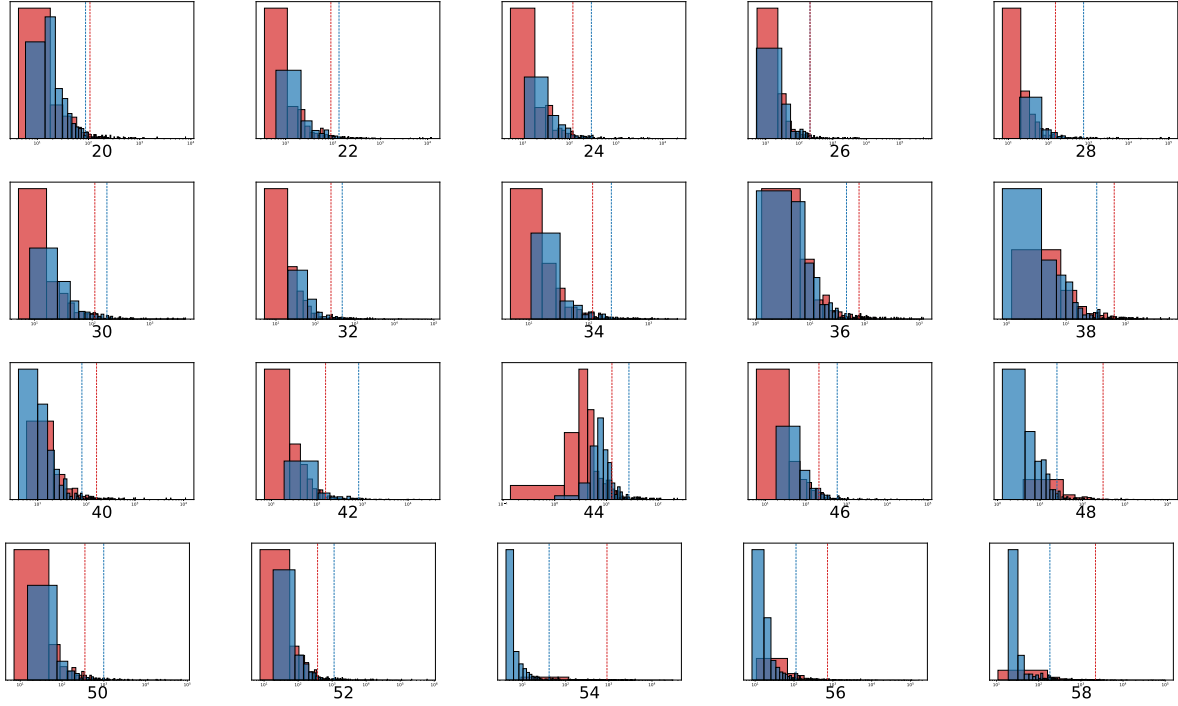


Figure 40: Histograms of the values of  $\|\nabla d_\omega/d_\omega\|_2$  and their 90th percentile across epochs. The epochs are displayed at the bottom of each histogram. The  $x$ -axis represents  $\|\nabla d_\omega/d_\omega\|_2$  values on a *logarithmic* scale, while the  $y$ -axis denotes density. Results are differentiated by color: **red** for the model with noise and **blue** for the model without noise. Preceding the 54th epoch where the noise-free GAN collapses, the noised model nearly always exhibits lower  $\|\nabla d_\omega/d_\omega\|_2$  values compared to its noise-free counterpart. Post 54th epoch, this relationship reverses. Notably, in the noise-free model,  $\|\nabla d_\omega/d_\omega\|_2$  tends towards zero, contributing to this observed divergence.

## G Rationale Behind the Discriminator Model

The rationale behind selecting the radius as  $2\sqrt{2}h$  emanates from the equivalence of an RBF kernel with bandwidth  $h$  to a Gaussian kernel of variance  $h^2/2$ . For a Gaussian distribution, 99 percent of samples lie within a sphere centered at the mean and extending to 3 times the standard deviation. We further extend the boundary to 4 times the standard deviation in light of potential mode mixture.

## H Extension to Other Divergence GANs

In this section, we outline how our framework can be extended to encompass other Divergence GANs. We focus on the  $f$ -GAN proposed in (Nowozin et al., 2016) with the  $f$ -divergence defined as

$$D_f(Q_\theta || p_{\text{data}}) = \int_{\mathbf{x}} p_{\text{data}}(\mathbf{x}) f\left(\frac{p_{\text{data}}(\mathbf{x})}{Q_\theta(\mathbf{x})}\right) d\mathbf{x}.$$

The variational lower bound of  $D_f(Q_\theta || p_{\text{data}})$  is used as the training objective:

$$F(\theta; \omega) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [g_f(V_\omega(\mathbf{x}))] + \mathbb{E}_{\mathbf{x} \sim Q_\theta} [-f^*(g_f(V_\omega(\mathbf{x})))].$$

Here,  $f^*$  is the Fenchel conjugate of  $f$ ,  $g_f$  is analogous to the generator and  $V_\omega$  is similar to the discriminator. We consider its variant where the objective function of the generator is modified to

$$-\mathbb{E}_{\mathbf{x} \sim Q_\theta} [g_f(V_\omega(\mathbf{x}))],$$

while the objective function of the discriminator remains unchanged.

**General methodology.** The key to analyzing Divergence GANs is their interpretation as particle models. The update of the generator  $Q_\theta$  can be recasted as:

- Generate particles  $Z_i = Q_\theta(z_i)$ .
- Update the particles  $\hat{Z}_i = Z_i + g'_f(V_\omega(Z_i))\nabla V_\omega(Z_i)$ .
- Update  $\theta$  by descending its stochastic gradient with respect to the Mean Square Error (MSE) loss between  $\hat{Z}_i$ 's and  $g(z_i)$ 's.

**Fitting phase.** We may plot the vector field  $g'_f(V_\omega(Z_i))\nabla V_\omega(Z_i)$  instead of the original  $\nabla d_\omega/d_\omega$  to visualize the updating process of particles, which promotes the fitting of the modes.

**Refining phase.** Only theorem 4.4 in section 4.2 needs to be modified to accommodate the desired Divergence GAN.

**Collapsing phase.** In section 5.1, apart from modifying the update formula for particles, a more appropriate model for the discriminator needs to be established in assumption 5.1 and a new threshold may be developed on the basis of it in algorithm 2.

## I Visualizing Generator Functions

This section discuss about how to plot fig. 3. While  $\Phi$  can be computed in MATLAB using the built-in function `normcdf`,  $\Psi^{-1}$  typically necessitates solving a non-linear equation at each evaluation point. To mitigate computational expenses, we choose to calculate the inverse of  $g$ , which is  $g^{-1} = \Phi^{-1} \circ \Psi$ . In this context,  $\Psi$  can be computed by employing `gmdistribution` to construct a Gaussian mixture model, followed by utilizing `cdf` to assess the cumulative distribution function (CDF) of the model at a specific point. To generate a plot of  $g$ , a mere interchange of the  $x$  and  $g^{-1}(x)$  in the `plot` function suffices.

## J Discussions

In this section, we provide additional intuitions and implications.

In terms of applicability scope, our theoretical framework is primarily derived from Divergence GANs, specifically NSGAN, where we can leverage their particle model interpretations. While Divergence GANs represent a significant category within GANs, they do not encompass some prominent GAN models, such as Wasserstein GAN with gradient penalty and MMD GAN. Exploring how our theoretical framework can be extended to incorporate these Integral Probability Metric (IPM) based GAN variants presents an intriguing avenue for future research.

Regarding the theoretical framework, it is important to note that not all Divergence GANs may fit neatly into the tripartite phases. While we often observe such empirical patterns, we acknowledge the possibility that when networks are not well-initialized or when advanced techniques are used, GAN training may deviate from the fitting phase entirely. However, these inquiries may spark independent interests and are beyond the scope of our study.

In our numerical experiments, we used relatively small-scale real-world datasets compared to modern datasets. This choice was deliberate as we aimed to assess the effectiveness of our early stopping algorithm in detecting the transition from refining to collapsing phases. Modern datasets often comprise exponentially more modes, which could potentially limit the efficacy of our algorithm, particularly considering that our algorithm takes the number of modes as an input parameter.