

Diffusion Models Learn Low-Dimensional Distributions via Subspace Clustering*

Peng Wang¹, Huijie Zhang¹, Zekai Zhang¹, Siyi Chen¹, Yi Ma², and Qing Qu¹

¹Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor

²Department of Electrical Engineering and Computer Science, University of California, Berkeley

December 27, 2024

Abstract

Recent empirical studies have demonstrated that diffusion models can effectively learn the image distribution and generate new samples. Remarkably, these models can achieve this even with a small number of training samples despite a large image dimension, circumventing the curse of dimensionality. In this work, we provide theoretical insights into this phenomenon by leveraging key empirical observations: (i) the low intrinsic dimensionality of image data, (ii) a union of manifold structure of image data, and (iii) the low-rank property of the denoising autoencoder in trained diffusion models. These observations motivate us to assume the underlying data distribution of image data as a mixture of low-rank Gaussians and to parameterize the denoising autoencoder as a low-rank model according to the score function of the assumed distribution. With these setups, we rigorously show that optimizing the training loss of diffusion models is equivalent to solving the canonical subspace clustering problem over the training samples. Based on this equivalence, we further show that the minimal number of samples required to learn the underlying distribution scales linearly with the intrinsic dimensions under the above data and model assumptions. This insight sheds light on why diffusion models can break the curse of dimensionality and exhibit the phase transition in learning distributions. Moreover, we empirically establish a correspondence between the subspaces and the semantic representations of image data, facilitating image editing. We validate these results with corroborated experimental results on both simulated distributions and image datasets. The code is available at <https://github.com/huijieZH/Diffusion-Model-Generalizability>.

Key words: Diffusion models, mixture of low-rank Gaussians, denoising autoencoder, phase transition

*The first and second authors contributed equally to this work. Correspondence to: Peng Wang (peng8wang@gmail.com).

Contents

1	Introduction	3
1.1	Our Contributions	3
1.2	Related Works	4
2	Problem Setup	7
2.1	Preliminaries on Score-Based Diffusion Models	7
2.2	Low-Dimensional Data and Models	8
3	Main Results	10
3.1	A Warm-Up Study: A Single Low-rank Gaussian Case	10
3.2	From Single Low-Rank Gaussian to Mixtures of Low-Rank Gaussians	12
4	Experiments & Practical Implications	13
4.1	Phase Transition in Learning Distributions	14
4.2	Semantic Meanings of Low-Dimensional Subspaces	16
5	Conclusion & Discussion	16
	Appendices	25
A	Proofs in Section 2	25
A.1	Relation between Score Matching Loss and Denoiser Autoencoder Loss	25
A.2	Proof of in Lemma 1	26
B	Proofs in Section 3.1	27
B.1	Proof of Theorem 1	27
B.2	Proof of Theorem 2	28
C	Proofs in Section 3.2	29
C.1	Theoretical Justification of the DAE (16)	29
C.2	Proof of Theorem 3	30
C.3	Proof of Theorem 4	31
D	Experimental Setups in Section 2.2	33
D.1	Verification of Mixture of Low-Rank Gaussian Data Distribution	34
D.2	Verification of Low-Rank Network Parameterization	35
E	Experimental Setups in Section 4	36
E.1	Learning the MoLRG distribution with the theoretical parameterization	36
E.2	Learning the MoLRG distribution with U-Net	37
E.3	Learning real-world image data distributions with U-Net	37
E.4	Correspondence between low-dimensional subspaces and image semantics	37
F	Auxiliary Results	38

1 Introduction

Generative modeling is a fundamental task in deep learning, which aims to learn a data distribution from training data to generate new samples. Recently, diffusion models have emerged as a powerful family of generative models, demonstrating remarkable performance across diverse domains, including image generation [1, 2, 3, 4], video content generation [5, 6], speech and audio synthesis [7, 8], fine-tuning [9, 10, 11] and solving inverse problem [12, 13, 14, 15]. In general, diffusion models learn a data distribution from training samples through a process that imitates the non-equilibrium thermodynamic diffusion process [2, 16, 17]. Specifically, the training and sampling of diffusion models involve two stages: (i) a forward diffusion process where Gaussian noise is incrementally added to training samples at each time step, and (ii) a backward sampling process where the noise is progressively removed through a neural network that is trained to approximate the score function at all time steps. As described in prior works [18, 17], the generative capability of diffusion models lies in their ability to learn the *score function* of the data distribution, i.e., the gradient of the logarithm of the probability density function (*pdf*). We refer the reader to [19, 20, 21] for a more comprehensive introduction and survey on diffusion models.

Despite the significant empirical success and recent advances in understanding sampling convergence [22, 23, 24], distribution approximation [25, 26], memorization [27, 28, 29, 30], and generalization [31, 32, 33] of diffusion models, the fundamental working mechanisms remain poorly understood. One of the key questions is

When and why do diffusion models learn the underlying data distribution without suffering from the curse of dimensionality?

In the existing literature, Oko et al. [26], Wibisono et al. [34] have demonstrated that learning an ϵ -accurate score via score estimation in the ℓ_2 -norm requires the number of training samples to grow at the rate of $O(\epsilon^{-n})$, where n is the ambient data dimension. This indicates that distribution learning via diffusion models is theoretically subject to the curse of dimensionality. However, recent empirical results in [31, 33] showed that the number of training samples for diffusion models to learn the underlying distribution is significantly *smaller* than the worst-case scenario, breaking the curse of dimensionality. Therefore, this discrepancy reveals a substantial gap between theory and practice.

In this work, we aim to address the above question of learning the underlying distribution via diffusion models by leveraging low-dimensional models. Our key observations are as follows: (i) The intrinsic dimensionality of real image data is significantly lower than the ambient dimension, a fact well-supported by extensive empirical evidence in [35, 36, 37]; (ii) Image data lies on a disjoint union of manifolds of varying intrinsic dimensions, as empirically verified in [38, 39, 40] (see Figure 1(a)); (iii) We empirically observe that the denoising autoencoder (DAE) [41, 42] of diffusion models trained on real-world image datasets exhibit low-rank structures (see Figure 3). Based on these observations, we conduct a theoretical investigation of distribution learning through diffusion models by assuming that (i) the underlying data distribution is a *mixture of low-rank Gaussians* (see Definition 1 and Figure 1(b)) and (ii) the denoising autoencoder is parameterized according to the score function of the MoLRG. Notably, these assumptions will be carefully discussed based on the existing literature and validated by our experiments on real image datasets.

1.1 Our Contributions

In this work, our contributions can be summarized as follows:

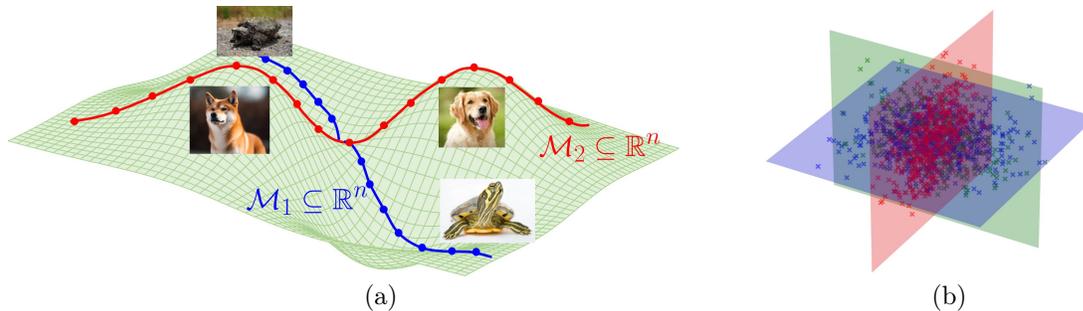


Figure 1: (a) Visualization of the union of manifold structure of image data. Here, different images lie on different manifolds $\mathcal{M}_i \subseteq \mathbb{R}^n$ of intrinsic dimension d with $d \ll n$. (b) An illustration of training samples that are generated according to the MoLRG model. This model is a local linearization of a union of manifolds.

- Equivalence between training diffusion models and subspace clustering.** Under the above setup, we show that the training loss of diffusion models based on the denoising autoencoder is equivalent to the *unsupervised* subspace clustering problem [43, 44, 45] (see Theorem 3). This equivalence implies that training diffusion models is essentially learning low-dimensional manifolds of the data distribution.
- Understanding phase transition from failure to success in learning distributions.** By leveraging the above equivalence and the data model, we further show that if the number of samples exceed the intrinsic dimension of the subspaces, the optimal solutions of the training loss can recover the underlying distribution. This explains why diffusion models can break the curse of dimensionality. Conversely, if the number of samples is insufficient, it may learn an incorrect distribution. This highlights a phase transition from failure to success in learning the underlying distributions as the sample size increases.
- Correspondence between semantic representations and the subspaces.** Interestingly, we find that the discovered low-dimensional subspaces in pre-trained diffusion models possess *semantic* meanings for natural images (see Figure 2). This motivates us to propose a training-free method to edit images on a frozen-trained diffusion model.

We also conduct extensive numerical experiments on both synthetic and real-world data sets to verify our assumptions and validate our theory. More broadly, the theoretical insights we gained in this work provide practical guidance as follows. First, we have shown that the number of samples for learning the underlying distribution via diffusion models scales proportionally with its intrinsic dimension. This insight allows us to improve training efficiency by quantifying the number of required training samples. Second, the identified correspondence between semantic representations and subspaces provides valuable guidance on controlling data generation. By manipulating the semantic representations within these subspaces, we can achieve more precise and targeted data generation.

1.2 Related Works

Now, we review recent works on diffusion models closely related to our study and discuss their connections to our work.

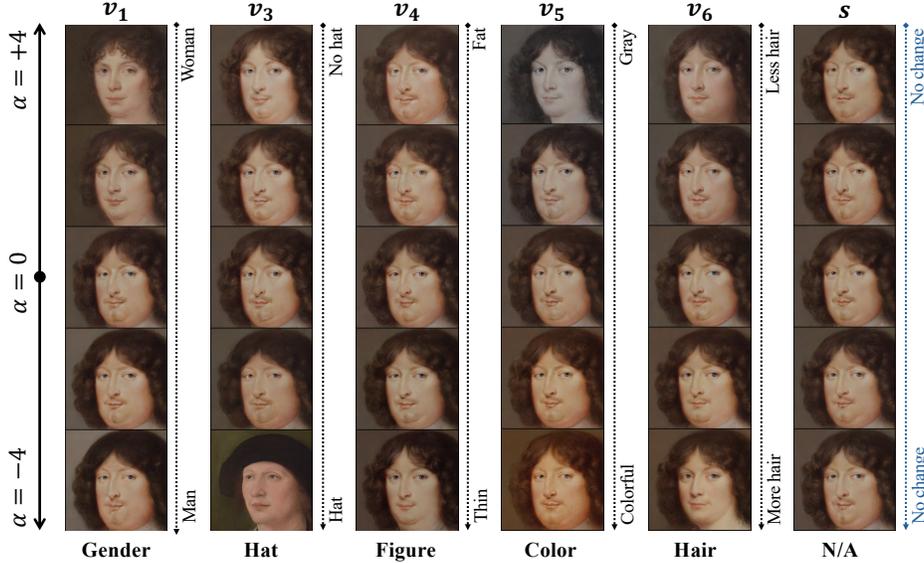


Figure 2: Correspondence between the singular vectors of the Jacobian of the DAE and semantic image attributes.

Learning a mixture of Gaussians via diffusion models. Recent works have extensively studied distribution learning and generalizability of diffusion models for learning a mixture of full-rank Gaussian (MoG) model [46, 47, 48, 49, 50]. Specifically, they assumed that there exist centers $\mu_1, \dots, \mu_K \in \mathbb{R}^n$ such that image data approximately follows from the following distribution:

$$\mathbf{x} \sim \sum_{k=1}^K \pi_k \mathcal{N}(\mu_k, \mathbf{I}_n), \quad (1)$$

where $\pi_k \geq 0$ is the mixing proportion of the k -th mixture component satisfying $\sum_{k=1}^K \pi_k = 1$. Shah et al. [49] showed that gradient descent for optimizing denoising diffusion probabilistic objective can efficiently locate the ground-truth centers of MoG. Later, Chen et al. [46] extended the above analysis without assuming the separation condition on the underlying mixture components. Gatmiry et al. [48] employed diffusion models to learn the MoG model and provided a time and sample complexity analysis to achieve a targeted total variation error. Recently, Wu et al. [50] provided a theoretical analysis of diffusion guidance in the MoG model. Moreover, Cole and Lu [47] analyzed the approximation and generalization of score-based diffusion models for learning sub-Gaussian distributions, which includes a mixture of Gaussians as an example.

Notably, the MoLRG model studied in our work is distinct from the above MoG model that is widely studied in the literature. Specifically, the MoG model consists of multiple Gaussians with varying means and covariance spanning the full-dimensional space (see Equation (1)), while a MoLRG comprises multiple Gaussians with zero mean and low-rank covariance (see Equation (7)), lying in a union of low-dimensional subspaces. As such, the MoLRG model, inspired by the inherent low-dimensionality of image datasets [35, 36, 37], offers a deeper insight into how diffusion models can learn underlying distributions in practice without suffering from the curse of dimensionality. Along the direction of studying diffusion models for low-dimensional data, Chen et al. [25] studied a general low-dimensional data distribution and showed that the underlying distribution can be recovered via diffusion models. In contrast, we specifically studied a mixture of low-dimensional

Gaussian distributions, enabling us to obtain a sharp bound for learning the underlying distribution.

Memorization and generalization in diffusion models. Recently, extensive studies [31, 32, 33] empirically revealed that diffusion models learn the score function across two distinct regimes — memorization (i.e., learning the empirical distribution) and generalization (i.e., learning the underlying distribution) — depending on the training dataset size vs. the model capacity. Notably, for a model with a fixed number of parameters, there is a phase transition from memorization to generalization as the number of training samples increases [31, 33]. In the memorization regime, recent works [27, 51] showed that diffusion models learn the empirical optimal denoiser when the network is highly over-parameterized. Yoon et al. [32] argued that diffusion models tend to generalize when they fail to memorize training data. In the generalization regime, extensive works [47, 52, 53] studied the generalization performance of diffusion models in different settings. Interestingly, Kadkhodaie et al. [31] observed that two distinct diffusion models can learn nearly identical score functions even if they are trained on sufficiently large separate, non-overlapping portions of the same dataset so that they can generate almost the same output from the same pure noise. Notably, most existing studies on the memorization and generalization of diffusion models are empirical. In contrast, our work provides rigorous theoretical explanations for these intriguing experimental observations based on the MoLRG model. We demonstrate that diffusion models learn the underlying data distribution with the number of training samples scaling linearly with the intrinsic dimension of the data distribution. Our theory reveals a phase transition from failure to success in learning the underlying distribution as the number of training samples increases, which sheds light on the phase transition from memorization to generalization.

Subspace clustering. Subspace clustering is a fundamental problem in unsupervised learning, which aims to identify and group data points that lie in a union of low-dimensional subspaces in a high-dimensional space [43, 44, 54]. Over the past years, a substantial body of literature has explored various approaches to the algorithmic development and theoretical analysis of subspace clustering. These include techniques such as sparse representation [55, 56, 57], low-rank representation [45, 58, 59], and spectral clustering [60, 61]. In this work, we present a new interpretation of diffusion models from the perspective of subspace clustering. This is the first time that diffusion models have been analyzed through this lens, offering new insights into how these models can effectively learn complex data distributions by leveraging the intrinsic low-dimensional subspaces within the data.

Notation. We write matrices in bold capital letters like \mathbf{A} , vectors in bold lower-case letters like \mathbf{a} , and scalars in plain letters like a . Given a matrix, we use $\|\mathbf{A}\|$ to denote its largest singular value (i.e., spectral norm), $\sigma_i(\mathbf{A})$ its i -th largest singular value, and a_{ij} its (i, j) -th entry, $\text{rank}(\mathbf{A})$ its rank, $\|\mathbf{A}\|_F$ its Frobenius norm. Given a vector \mathbf{a} , we use $\|\mathbf{a}\|$ to denote its Euclidean norm and a_i its i -th entry. Let $\mathcal{O}^{n \times d}$ denote the set of all $n \times d$ orthonormal matrices. We simply write the score function $\nabla_{\mathbf{x}} \log p(\mathbf{x})$ of a distribution with probability density function (pdf) $p(\mathbf{x})$ as $\nabla \log p(\mathbf{x})$.

Organization. In Section 2, we introduce the preliminaries on diffusion models and some assumptions on the data and model. In Section 3, we present the main results of this paper. In Section 4, we conduct numerical experiments to support our theory and demonstrate its practical implications. Finally, in Section 5, we summarize our work and discuss potential directions for future research.

2 Problem Setup

In this work, we consider an image dataset consisting of samples $\{\mathbf{x}^{(i)}\}_{i=1}^N \subseteq \mathbb{R}^n$, where each data point is *i.i.d.* sampled from an underlying data distribution $p_{\text{data}}(\mathbf{x})$. Instead of learning this pdf directly, score-based diffusion models aim to learn the score of this distribution from the training samples.

2.1 Preliminaries on Score-Based Diffusion Models

Forward and reverse SDEs of diffusion models. In general, diffusion models consist of forward and reverse processes indexed by a continuous time variable $t \in [0, 1]$. Specifically, the forward process progressively injects noise into the data. This process can be described by the following stochastic differential equation (SDE):

$$d\mathbf{x}_t = f(t)\mathbf{x}_t dt + g(t)d\mathbf{w}_t, \quad (2)$$

where $\mathbf{x}_0 \sim p_{\text{data}}$, the scalar functions $f(t), g(t) : \mathbb{R} \rightarrow \mathbb{R}$ respectively denote the drift and diffusion coefficients,¹ and $\{\mathbf{w}_t\}_{t \in [0,1]}$ is the standard Wiener process. For ease of exposition, let $p_t(\mathbf{x})$ denote the pdf of \mathbf{x}_t and $p_t(\mathbf{x}_t|\mathbf{x}_0)$ the transition kernel from \mathbf{x}_0 to \mathbf{x}_t . According to (2), we have

$$p_t(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; s_t\mathbf{x}_0, s_t^2\sigma_t^2\mathbf{I}_n), \text{ where } s_t = \exp\left(\int_0^t f(\xi)d\xi\right), \sigma_t = \sqrt{\int_0^t \frac{g^2(\xi)}{s^2(\xi)}d\xi}. \quad (3)$$

The reverse process gradually removes the noise from \mathbf{x}_1 via the following reverse-time SDE:

$$d\mathbf{x}_t = (f(t)\mathbf{x}_t - g^2(t)\nabla \log p_t(\mathbf{x}_t)) dt + g(t)d\bar{\mathbf{w}}_t, \quad (4)$$

where $\{\bar{\mathbf{w}}_t\}_{t \in [0,1]}$ is another standard Wiener process, independent of $\{\mathbf{w}_t\}$, running backward in time from $t = 1$ to $t = 0$. It is worth noting that if \mathbf{x}_1 and $\nabla \log p_t$ are provided, the reverse process has exactly the same distribution as the forward process at each time $t \geq 0$ [63].

Training loss of diffusion models. Unfortunately, the score function $\nabla \log p_t$ is usually unknown, as it depends on the underlying data distribution p_{data} . To enable data generation via the reverse SDE (4), a common approach is to estimate the score function $\nabla \log p_t$ using the training samples $\{\mathbf{x}^{(i)}\}_{i=1}^N$ based on the scoring matching [2, 17]. Because of the equivalence between the score function $\nabla \log p_t(\mathbf{x}_t)$ and the posterior mean $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$, i.e.,

$$s_t \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t] = \mathbf{x}_t + s_t^2 \sigma_t^2 \nabla \log p_t(\mathbf{x}_t), \quad (5)$$

according to Tweedie’s formula and (3), an alternative approach to estimating the score function $\nabla \log p_t$ is to estimate the posterior mean $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$. Consequently, extensive works [64, 31, 62, 42, 65] have considered training a time-dependent function $\mathbf{x}_\theta(\cdot, t) : \mathbb{R}^n \times [0, 1] \rightarrow \mathbb{R}^n$, known as *denoising autoencoder* (DAE), parameterized by a neural network with parameters θ to estimate the posterior mean $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$. To determine the parameters θ , we can minimize the following empirical loss:

$$\min_{\theta} \ell(\theta) := \frac{1}{N} \sum_{i=1}^N \int_0^1 \lambda_t \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)} \left[\left\| \mathbf{x}_\theta(s_t \mathbf{x}^{(i)} + \gamma_t \epsilon, t) - \mathbf{x}^{(i)} \right\|^2 \right] dt, \quad (6)$$

¹In general, the functions $f(t)$ and $g(t)$ are chosen such that (i) \mathbf{x}_t for all t close to 0 approximately follows the data distribution p_{data} and (ii) \mathbf{x}_t for all t close to 1 is nearly a standard Gaussian distribution; see, e.g., the settings in [2, 62, 17].

where $\lambda_t : [0, 1] \rightarrow \mathbb{R}^+$ is a weighting function and $\gamma_t := s_t \sigma_t$. As shown in [42], training the DAE is equivalent to performing explicit or implicit score matching under mild conditions. We refer the reader to Section A.1 for the relationship between this loss and the score-matching loss in [17, 42].

2.2 Low-Dimensional Data and Models

Mixture of low-rank Gaussian data distribution. Although real-world image datasets are high dimensional in terms of pixel count and overall data volume, extensive empirical works [35, 39, 36, 37] suggest that their intrinsic dimensions are much lower. For instance, Pope et al. [36] employed a kernelized nearest neighbor method to estimate the intrinsic dimensionality of various datasets, including MNIST [66], CIFAR-10 [67], and ImageNet [68]. Their findings indicate that even for complex datasets like ImageNet, the intrinsic dimensionality is approximately 40, which is significantly lower than its ambient dimension.

Recently, Brown et al. [38], Kamkari et al. [39] empirically validated the *union of manifolds* hypothesis, demonstrating that high-dimensional image data often lies on a disjoint union of manifolds instead of a single manifold. Notably, the union of subspaces, a special case of this hypothesis, has been widely used in the subspace clustering literature [55, 44, 45]. These observations motivate us to model the underlying data distribution as a *mixture of low-rank Gaussians*, where the data points are generated from a mixture of several Gaussian distributions with zero mean and low-rank covariance matrices. We formally define the MoLRG distribution as follows:

Definition 1 (Mixtures of Low-Rank Gaussians). *We say that a random vector $\mathbf{x} \in \mathbb{R}^n$ follows a mixture of K low-rank Gaussian distribution with parameters $\{\pi_k\}_{k=1}^K$ and $\{\mathbf{U}_k^*\}_{k=1}^K$ if*

$$\mathbf{x} \sim \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{0}, \mathbf{U}_k^* \mathbf{U}_k^{*T}), \quad (7)$$

where $\mathbf{U}_k^* \in \mathcal{O}^{n \times d_k}$ denotes the orthonormal basis of the k -th component and $\pi_k \geq 0$ is the mixing proportion of the k -th mixture component satisfying $\sum_{k=1}^K \pi_k = 1$.

Before we proceed, we make some remarks on this data model. First, to study how diffusion models learn the underlying data distribution, many recent works have studied a mixture of full-rank Gaussian distributions (see Eq. (1)); see, e.g., [46, 48, 49, 69, 50]. However, compared to this model, a MoLRG is a more suitable model for capturing the low-dimensionality in image data distribution. Second, Brown et al. [38], Kamkari et al. [39] conducted extensive numerical experiments to validate that image datasets such as MNIST and ImageNet approximately lie on a union of low-dimensional manifolds. Because a nonlinear manifold can be well approximated by its tangent space (i.e., a linear subspace) in a local neighborhood, the MoLRG model, which represents data as a union of linear subspace, serves a good local approximation of a union of manifolds. Finally, assuming Gaussian distributions in each subspace in the MoLRG model is to ensure theoretical tractability while approximating the real-world image distributions, making it a practical starting point for theoretical studies on real-world image datasets. Now, we compute the ground-truth posterior mean $\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]$ when \mathbf{x}_0 satisfies the MoLRG model as follows.

Lemma 1. *Suppose that \mathbf{x}_0 satisfies the MoLRG model. For each time $t > 0$, it holds that*

$$\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = \frac{s_t}{s_t^2 + \gamma_t^2} \frac{\sum_{k=1}^K \pi_k \exp(\phi_t \|\mathbf{U}_k^{*T} \mathbf{x}_t\|^2) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}_t}{\sum_{k=1}^K \pi_k \exp(\phi_t \|\mathbf{U}_k^{*T} \mathbf{x}_t\|^2)}, \quad \text{where } \phi_t := \frac{s_t^2}{2\gamma_t^2(s_t^2 + \gamma_t^2)}. \quad (8)$$

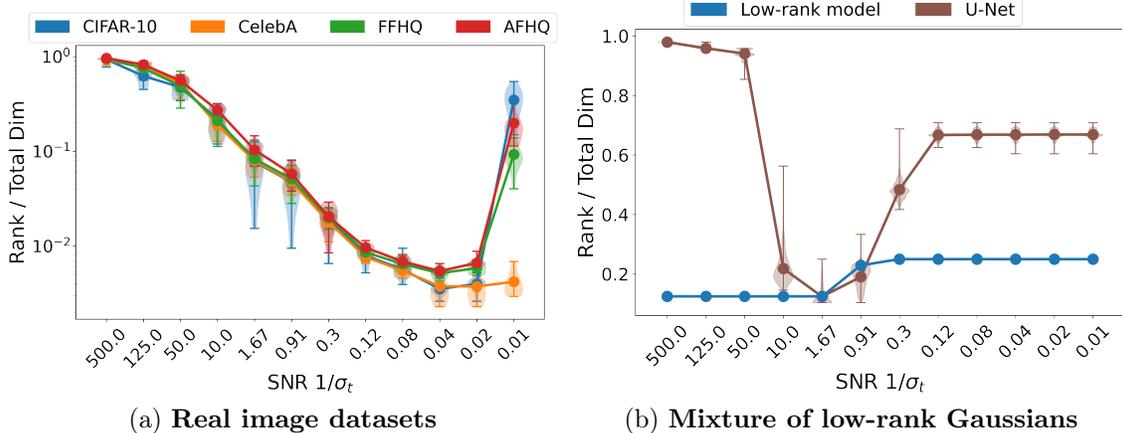


Figure 3: **Low-rank property of the denoising autoencoder of trained diffusion models.** We plot the ratio of the numerical rank of the Jacobian of the denoising autoencoder, i.e., $\nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t)$, over the total dimension against the signal-to-noise ratio (SNR) $1/\sigma_t$ by training diffusion models on different datasets. (a) We train diffusion models on image datasets CIFAR-10 [67], CelebA [70], FFHQ [71], and AFHQ [72]. The experimental details are provided in Section D.1. (b) We respectively train diffusion models with the low-rank parameterization (9) and U-Net on a mixture of low-rank Gaussian distributions. The experimental details are provided in Section D.2.

We defer the proof of this lemma to Section A.2. Notably, this lemma shows that the ground-truth posterior mean is a convex combination of projections of \mathbf{x}_t onto each subspace, where the weights are soft-max functions. Moreover, this lemma provides guidance on the network parameterization of the DAE $\mathbf{x}_\theta(\cdot, t)$ as discussed below.

Low-rank network parameterization. In this work, we empirically observed that the DAE $\mathbf{x}_\theta(\cdot, t)$ trained on real-world image datasets exhibits a low-dimensional structure. Specifically, when we train diffusion models with the U-Net architecture [73] on various image datasets, it is observed that the numerical rank of the Jacobian of the DAE, i.e., $\nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t)$, is substantially lower than the ambient dimension in most time steps; see Figure 3(a). Additionally, this pattern of low dimensionality appears to be consistent across different datasets with different noise levels t .

When training diffusion models with U-Net on the samples generated according to the MoLRG model, the Jacobian of the DAE also exhibits a similar low-rank pattern, as illustrated in Figure 3(b). For the theoretical study based upon MoLRG, the above observations motivate us to consider a low-rank parameterization of the network. According to the ground-truth posterior mean of the MoLRG model in Lemma 1, a natural parameterization for the DAE is

$$\mathbf{x}_\theta(\mathbf{x}_t, t) = \frac{s_t}{s_t^2 + \gamma_t^2} \sum_{k=1}^K w_k(\boldsymbol{\theta}; \mathbf{x}_t) \mathbf{U}_k \mathbf{U}_k^T \mathbf{x}_t, \text{ where } w_k(\boldsymbol{\theta}; \mathbf{x}_t) = \frac{\pi_k \exp(\phi_t \|\mathbf{U}_k^T \mathbf{x}_t\|^2)}{\sum_{l=1}^K \pi_l \exp(\phi_t \|\mathbf{U}_l^T \mathbf{x}_t\|^2)} \quad (9)$$

and the network parameters $\boldsymbol{\theta} = \{\mathbf{U}_k\}_{k=1}^K$ satisfy $\mathbf{U}_k \in \mathcal{O}^{n \times d_k}$. Although this approach may seem idealized, it offers several practical insights:

- **Resemblance to the U-Net architecture:** If we consider a single low-rank Gaussian, the network parameterization for approximating the score function takes the form $\mathbf{x} - s_t/(s_t^2 +$

$\gamma_t^2)UU^T\mathbf{x}$, which resembles the structure of a practical U-Net with a linear encoder, decoder, and skip connections. This provides theoretical insights into why U-Net is preferred for training diffusion models.

- **Capturing low-rank property of the DAE:** As shown in Figure 3, the Jacobian of the DAE of diffusion models trained on real image data is low-rank for most time steps, and our parameterization captures this low-rank property.
- **Providing practical guidance:** The theoretical insights gained from this work offer practical guidance. First, to learn the underlying distribution, the number of samples should be proportional to its intrinsic dimension. In practice, this informs us on how to use a minimal number of samples to train diffusion models to achieve generalization. Second, in Section 4.2, we empirically showed that the subspace basis has semantic meanings. Based upon this, our results shed light on how to control the data generation by manipulating the semantic representation within the subspaces.

Similar simplifications have been widely used for theoretical analysis in various ideal data distributions; see, e.g., [25, 46, 48, 49]. Notably, under this specific network parameterization in (9), learning the score function $\nabla \log p_t(\mathbf{x}_t)$ reduces to learning the network parameters $\boldsymbol{\theta}$ in (9) according to Lemma 1 and Eq. (5).

3 Main Results

Based on the setups in Section 2.2, we are ready to conduct a theoretical analysis of distribution learning using diffusion models. Specifically, when the underlying data distribution $p_{\text{data}}(\mathbf{x})$ is modeled as MoLRG and the DAE is approximately parameterized as Eq. (9), we show that

- The training loss (6) of diffusion models is *equivalent* to the **subspace clustering** problem, which can be reduced to the PCA problem in a single low-rank Gaussian case.
- The MoLRG distribution can be learned by optimizing the training loss when the minimum number of samples scales linearly with the **intrinsic dimensionality** of the data.

Note that $\min_{\boldsymbol{\theta} \in \Theta} \ell(\boldsymbol{\theta})$ is equivalent to $\min_{\boldsymbol{\theta} \in \Theta} f(\boldsymbol{\theta})$ when we have $f(\boldsymbol{\theta}) = a\ell(\boldsymbol{\theta}) + b$, where $a > 0$ and b are absolute constants that do not depend on $\boldsymbol{\theta}$. This equivalence applies to our results in both Theorem 1 and Theorem 3.

3.1 A Warm-Up Study: A Single Low-rank Gaussian Case

To begin, we start from a simple case that the underlying distribution p_{data} is a *single* low-rank Gaussian. Specifically, the training samples $\{\mathbf{x}^{(i)}\}_{i=1}^N \subseteq \mathbb{R}^n$ are generated according to

$$\mathbf{x}^{(i)} = U^* \mathbf{a}_i + \mathbf{e}_i, \tag{10}$$

where $\mathbf{U}^* \in \mathcal{O}^{n \times d}$ denotes an orthonormal basis, $\mathbf{a}_i \stackrel{i.i.d.}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ is coefficients for each $i \in [N]$, and $\mathbf{e}_i \in \mathbb{R}^n$ is noise for all $i \in [N]$.² According to (9), we parameterize the DAE into

$$\mathbf{x}_\theta(\mathbf{x}_t, t) = \frac{s_t}{s_t^2 + \gamma_t^2} \mathbf{U} \mathbf{U}^T \mathbf{x}_t, \quad (11)$$

where $\theta = \mathbf{U} \in \mathcal{O}^{n \times d}$. Equipped with the above setup, we can show the following result.

Theorem 1. *Suppose that the DAE $\mathbf{x}_\theta(\cdot, t)$ in Problem (6) is parameterized into (11) for each $t \in [0, 1]$. Then, Problem (6) is equivalent to the following principal component analysis (PCA) problem:*

$$\max_{\mathbf{U} \in \mathbb{R}^{n \times d}} \sum_{i=1}^N \|\mathbf{U}^T \mathbf{x}^{(i)}\|^2 \quad \text{s.t.} \quad \mathbf{U}^T \mathbf{U} = \mathbf{I}_d. \quad (12)$$

We defer the proof to Section B.1. In the single low-rank Gaussian model, Theorem 1 demonstrates that training diffusion models with a DAE of the form (11) to learn this distribution is equivalent to performing PCA on the training samples. Leveraging this equivalence, we can further characterize the number of samples required for learning underlying distribution under the data model (10).

Theorem 2. *Consider the setting of Theorem 1. Suppose that the training samples $\{\mathbf{x}^{(i)}\}_{i=1}^N$ are generated according to the noisy single low-rank Gaussian model defined in (10). Let $\hat{\mathbf{U}}$ denote an optimal solution of Problem (6). The following statements hold:*

- i) *If $N \geq d$, it holds with probability at least $1 - 1/2^{N-d+1} - \exp(-c_2 N)$ that any optimal solution $\hat{\mathbf{U}}$ satisfies*

$$\left\| \hat{\mathbf{U}} \hat{\mathbf{U}}^T - \mathbf{U}^* \mathbf{U}^{*T} \right\|_F \leq \frac{c_1 \sqrt{\sum_{i=1}^N \|\mathbf{e}_i\|^2}}{\sqrt{N} - \sqrt{d-1}}, \quad (13)$$

where $c_1, c_2 > 0$ are constants that depend polynomially only on the Gaussian moment.

- ii) *If $N < d$, there exists an optimal solution $\hat{\mathbf{U}} \in \mathcal{O}^{n \times d}$ such that with probability at least $1 - 1/2^{d-N+1} - \exp(-c'_2 d)$,*

$$\left\| \hat{\mathbf{U}} \hat{\mathbf{U}}^T - \mathbf{U}^* \mathbf{U}^{*T} \right\|_F \geq \sqrt{2 \min\{d - N, n - d\}} - \frac{c'_1 \sqrt{\sum_{i=1}^N \|\mathbf{e}_i\|^2}}{\sqrt{d} - \sqrt{N-1}}, \quad (14)$$

where $c'_1, c'_2 > 0$ are constants that depend polynomially only on the Gaussian moment.

Remark 1. We defer the proof to Section B.2. Building on the equivalence in Theorem 1 and the DAE parameterization (11), Theorem 2 clearly shows a phase transition from failure to success in learning the underlying distribution as the number of training samples increases. This phase transition is further corroborated by our experiments in Figures 4(a) and 4(b).

- When the number of training samples is larger than the dimension of the subspace, i.e., $N \geq d$, any optimal solution $\hat{\mathbf{U}}$ recovers the basis of underlying subspace with an approximation error depending on the noise level. This, together with the fact that learning the underlying distributions corresponds to learning the parameters \mathbf{U} under the parameterization (11), implies that optimizing the training loss of diffusion models can learn the underlying distribution. Conversely, when $N < d$, optimizing the training loss fails to learn the underlying distribution.

²Since real-world images inherently contain noise due to various factors, such as sensor limitation, environment conditions, and transition error, it is reasonable to add a noise term to this model.

- Note that our theory cannot explain why diffusion models memorize training data (i.e., learning the empirical distribution). This is because the parameterization (11) is not as sufficiently over-parameterized as architectures like U-Net. We plan to explore this over-parameterized setting in future work to better understand how diffusion models achieve memorization and to extend our theoretical insights accordingly.

3.2 From Single Low-Rank Gaussian to Mixtures of Low-Rank Gaussians

In this subsection, we extend the above study to the MoLRG distribution. In particular, we consider a noisy version of the MoLRG model as defined Definition 1. Specifically, the training samples are generated by

$$\mathbf{x}^{(i)} = \mathbf{U}_k^* \mathbf{a}_i + \mathbf{e}_i \text{ with probability } \pi_k, \quad \forall i \in [N], \quad (15)$$

where $\mathbf{U}_k^* \in \mathcal{O}^{n \times d_k}$ denotes an orthonormal basis for each $k \in [K]$, $\mathbf{a}_i \stackrel{i.i.d.}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_{d_k})$ is coefficients, and $\mathbf{e}_i \in \mathbb{R}^n$ is noise for each $i \in [N]$. As argued by [38], image data lies on a *disjoint* union of manifolds. This motivates us to assume that the basis matrices of subspaces satisfy $\mathbf{U}_k^{*T} \mathbf{U}_l^* = \mathbf{0}$ for each $k \neq l$. To simplify our analysis, we assume that $d_1 = \dots = d_K = d$ and the mixing weights satisfy $\pi_1 = \dots = \pi_K = 1/K$. Moreover, we consider a hard-max counterpart of Eq. (9) for the DAE parameterization as follows:

$$\mathbf{x}_\theta(\mathbf{x}_t, t) = \frac{s_t}{s_t^2 + \gamma_t^2} \sum_{k=1}^K \hat{w}_k(\theta, \mathbf{x}_0) \mathbf{U}_k \mathbf{U}_k^T \mathbf{x}_t, \quad (16)$$

where $\theta = \{\mathbf{U}_k\}_{k=1}^K$ and the weights $\{\hat{w}_k(\theta; \mathbf{x}_0)\}_{k=1}^K$ are set as

$$\hat{w}_k(\theta; \mathbf{x}_0) = \begin{cases} 1, & \text{if } k = k_0, \\ 0, & \text{otherwise,} \end{cases} \quad (17)$$

where $k_0 \in [K]$ is an index satisfying $\|\mathbf{U}_{k_0}^T \mathbf{x}_0\| \geq \|\mathbf{U}_l^T \mathbf{x}_0\|$ for all $l \neq k_0 \in [K]$. We should point out that we use two key approximations here. First, the soft-max weights $\{w_k(\theta, \mathbf{x}_t)\}$ in Eq. (9) are approximated by the hard-max weights $\{\hat{w}_k(\theta; \mathbf{x}_0)\}_{k=1}^K$. Second, $\|\mathbf{U}_k^T \mathbf{x}_t\|$ is approximated by its expectation, i.e., $\mathbb{E}_\epsilon[\|\mathbf{U}_k^T \mathbf{x}_t\|^2] = \mathbb{E}_\epsilon[\|\mathbf{U}_k^T (s_t \mathbf{x}_0 + \gamma_t \epsilon)\|^2] = s_t^2 \|\mathbf{U}_k^T \mathbf{x}_0\|^2 + \gamma_t^2 d$. We refer the reader to Section C.1 for more details on these approximation. Now, we are ready to show the following theorem.

Theorem 3. *Suppose that the DAE $\mathbf{x}_\theta(\cdot, t)$ in Problem (6) is parameterized into (16) for each $t \in [0, 1]$, where $\hat{w}_k(\theta, \mathbf{x}_0)$ is defined in (17) for each $k \in [K]$. Then, Problem (6) is equivalent to the following subspace clustering problem:*

$$\max_{\theta} \frac{1}{N} \sum_{k=1}^K \sum_{i \in C_k(\theta)} \|\mathbf{U}_k^T \mathbf{x}^{(i)}\|^2 \quad \text{s.t.} \quad [\mathbf{U}_1, \dots, \mathbf{U}_K] \in \mathcal{O}^{n \times dK}, \quad (18)$$

where $C_k(\theta) := \{i \in [N] : \|\mathbf{U}_k^T \mathbf{x}^{(i)}\| \geq \|\mathbf{U}_l^T \mathbf{x}^{(i)}\|, \forall l \neq k\}$ for each $k \in [K]$.

We defer the proof to Section C.2. When the DAE is parameterized into (16), Theorem 3 demonstrates that optimizing the training loss of diffusion models is equivalent to solving the subspace clustering problem [43, 44, 45]. Problem (18) seeks to maximize the sum of the squared norms of each point's projection onto its assigned subspace, subject to the constraint $C_k(\theta)$. Moreover, the equivalence enables us to characterize the required minimum number of samples for learning the underlying MoLRG distribution.

Theorem 4. Consider the setting of Theorem 3. Suppose that the training samples $\{\mathbf{x}^{(i)}\}_{i=1}^N$ are generated by the MoLRG distribution in Definition 1. Suppose $d \gtrsim \log N$ and $\|\mathbf{e}_i\| \lesssim \sqrt{d/N}$ for all $i \in [N]$. Let $\{\hat{\mathbf{U}}_k\}_{k=1}^K$ denote an optimal solution of Problem (6) and N_k denote the number of samples from the k -th Gaussian component. Then, the following statements hold:

- (i) If $N_k \geq d$ for each $k \in [K]$, there exists a permutation $\Pi : [K] \rightarrow [K]$ such that with probability at least $1 - 2K^2N^{-1} - \sum_{k=1}^K (1/2^{N_k-d+1} + \exp(-c_2N_k))$ for each $k \in [K]$,

$$\left\| \hat{\mathbf{U}}_{\Pi(k)} \hat{\mathbf{U}}_{\Pi(k)}^T - \mathbf{U}_k^* \mathbf{U}_k^{*T} \right\|_F \leq \frac{c_1 \sqrt{\sum_{i=1}^N \|\mathbf{e}_i\|^2}}{\sqrt{N_k} - \sqrt{d-1}}, \quad (19)$$

where $c_1, c_2 > 0$ are constants that depend polynomially only on the Gaussian moment.

- (ii) If $N_k < d$ for some $k \in [K]$, there exists a permutation $\Pi : [K] \rightarrow [K]$ and $k \in [K]$ such that with probability at least $1 - 2K^2N^{-1} - \sum_{k=1}^K (1/2^{d-N_k+1} + \exp(-c'_2N_k))$,

$$\left\| \hat{\mathbf{U}}_{\Pi(k)} \hat{\mathbf{U}}_{\Pi(k)}^T - \mathbf{U}_k^* \mathbf{U}_k^{*T} \right\|_F \geq \sqrt{2 \min\{d - N_k, n - d\}} - \frac{c'_1 \sqrt{\sum_{i=1}^N \|\mathbf{e}_i\|^2}}{\sqrt{d} - \sqrt{N_k - 1}}, \quad (20)$$

where $c'_1, c'_2 > 0$ are constants that depend polynomially only on the Gaussian

Remark 2. We defer the proof to Section C.3. We discuss the implications of our results below.

- *Phase transition in learning the underlying distribution.* This theorem demonstrates that when the number of samples in each subspace exceeds the dimension of the subspace and the noise is bounded, the optimal solution of the training loss (6) under the parameterization (16) can recover the underlying subspaces up to the noise level. Conversely, when the number of samples is insufficient, there exists an optimal solution that may recover wrong subspaces; see Figures 4(c) and 4(d).
- *Connections to the phase transition from memorization to generalization.* We should clarify the difference between the phase transition described in Theorems 2 & 4 and the phase transition from memorization to generalization. Our phase transition refers to the shift from failure to success of learning the underlying distribution as the number of training samples increase, whereas the latter concerns the shift from memorizing data to generalizing from it as the number of training samples increases. Nevertheless, our theory still sheds light on the minimal number of samples required for diffusion models to enter the generalized regime.
- *Semantic meanings of the subspaces.* More interestingly, we empirically discovered a correspondence between the subspace bases and the semantics of images, as demonstrated by the experiments in Section 4.2. Specifically, when we move \mathbf{x}_t along the direction of a singular vector of the Jacobian of the DAE trained on real image data with step size, the corresponding semantics of the generated images changes proportionally to the step size; see Figures 6 and 9(a, c). This enables controlled image generation without the need for additional training.

4 Experiments & Practical Implications

In this section, we first investigate phase transitions of diffusion models in learning distributions under both theoretical and practical settings in Section 4.1. Next, we demonstrate the practical

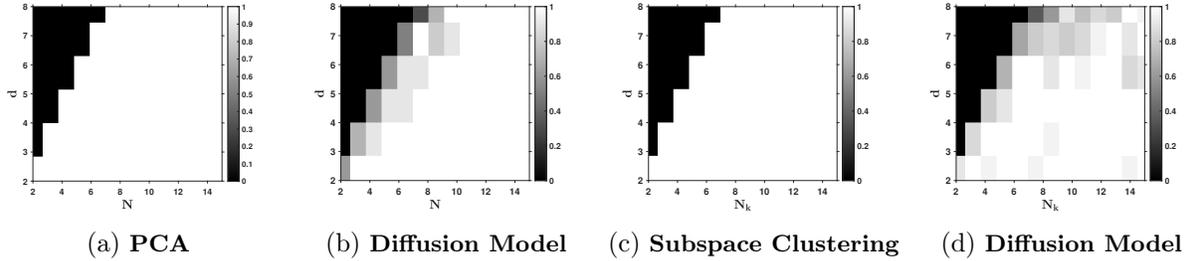


Figure 4: **Phase transition of learning the MoLRG distribution.** The x -axis is the number of training samples and y -axis is the dimension of subspaces. Darker pixels represent a lower empirical probability of success. When $K = 1$, we apply SVD and train diffusion models to solve Problems (12) and (6), visualizing the results in (a) and (b), respectively. When $K = 2$, we apply a subspace clustering method and train diffusion models for solving Problems (18) and (6), visualizing the results in (c) and (d), respectively. Additional experiments for the case when $K = 3$ are presented in Figure 8.

implications of our work by exploring the correspondence between low-dimensional subspaces and semantic representations for controllable image editing in Section 4.2. Detailed experimental setups are postponed to Section E.

4.1 Phase Transition in Learning Distributions

In this subsection, we conduct experiments on both synthetic and real datasets to study the phase transition of diffusion models in learning distributions.

Learning the MoLRG distribution with the theoretical parameterizations. To begin, we optimize the training loss (6) with the theoretical parameterization (9), where the data samples are generated by the MoLRG distribution. First, we apply stochastic gradient descent (see Algorithm 1) to solve Problem (6) with the DAE parameterized as (9). For comparison, according to Theorem 1 (resp., Theorem 3), we apply a singular value decomposition (resp., subspace clustering [45]) to solve Problem (12) (resp, Problem (18)).

We conduct three sets of experiments, where the data samples are respectively generated according to the single low-rank Gaussian distribution (10) with $K = 1$ and a mixture of low-rank Gaussian distributions (15) with $K = 2, 3$. In each set, we set the total dimension $n = 48$ and let the subspace dimension d and the number of training samples N vary from 2 to 8 and 2 to 15 with increments of 1, respectively. For every pair of d and N , we generate 20 instances, run the above methods, and calculate the successful rate of recovering the underlying subspaces. Then, the simulation results are visualized in Figure 4 and Figure 8. It is observed that all these methods exhibit a phase transition from failure to success in learning the subspaces as the number of training samples increases, which supports the results in Theorems 2 and 4.

Learning the MoLRG distribution with U-Net. Next, we optimize the training loss (6) with parameterizing the DAE $\mathbf{x}_\theta(\cdot, t)$ using U-Net, detailed experiment settings are in Section E.2. Since U-Net is highly over-parameterized with about 54 million parameters and is different from the posterior mean in Lemma 1, we measure the generalization ability of U-Net via *generalization (GL)*

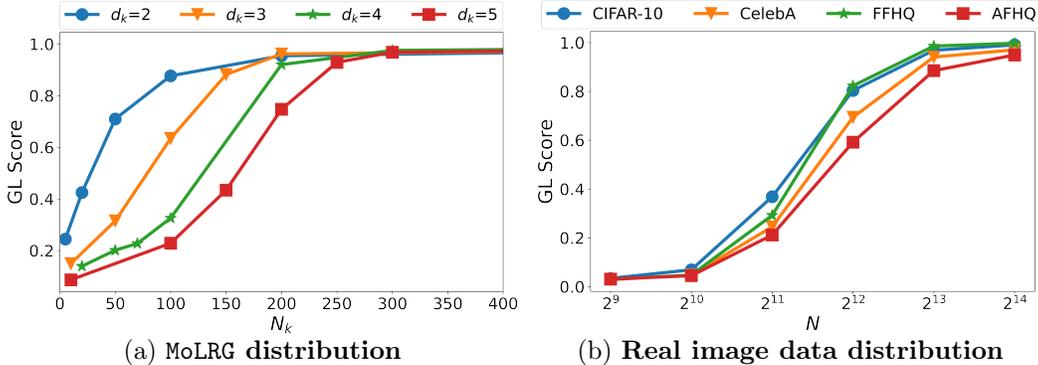


Figure 5: **Phase transition of learning distributions via U-Net.** In (a), the x -axis is the number of training samples over the intrinsic dimension, while in (b), it is the total number of training samples. The y -axis is the GL score. We train diffusion models with the U-Net architecture on (a) the data samples generated by the MoLRG distribution with $K = 2$, $n = 48$ and d_k varying from 3 to 6 and (b) real image datasets CIFAR-10, CelebA, FFHQ and AFHQ. The GL score is low when U-Net memorizes the training data and high when it learns the underlying distribution.

score [33] defined as follows:

$$\text{GL score} = \frac{\mathcal{D}(\mathbf{x}_{\text{gen}}^{(i)})}{\mathcal{D}(\mathbf{x}_{\text{MoLRG}}^{(i)})}, \quad \mathcal{D}(\mathbf{x}^{(i)}) := \sum_{j=1}^N \min_{j \neq i} \|\mathbf{x}^{(i)} - \mathbf{x}^{(j)}\|, \quad (21)$$

where $\{\mathbf{x}_{\text{MoLRG}}^{(i)}\}_{i=1}^N$ are samples generated from the MoLRG distribution and $\{\mathbf{x}_{\text{gen}}^{(i)}\}_{i=1}^N$ are new samples generated by the trained U-Net. Intuitively, $\mathcal{D}(\mathbf{x}_{\text{gen}}^{(i)})$ reflects the uniformity of samples in the space: its value is small when the generated samples cluster around the training data, while the value is large when generated samples disperse in the entire space. Therefore, the trained diffusion model is in memorization regime when $\mathcal{D}(\mathbf{x}_{\text{gen}}^{(i)}) \ll \mathcal{D}(\mathbf{x}_{\text{MoLRG}}^{(i)})$ and the GL score is close to 0, while it is in generalization regime when $\mathcal{D}(\mathbf{x}_{\text{gen}}^{(i)}) \approx \mathcal{D}(\mathbf{x}_{\text{MoLRG}}^{(i)})$ and the GL score is close to 1.

In the experiments, we generate the data samples using the MoLRG distribution with $K = 2$, $n = 48$, and $d_k \in \{3, 4, 5, 6\}$. Then, we plot the GL score against the N_k/d_k for each d_k in Figure 5(a). It is observed that for a fixed d_k , the generalization performance of diffusion models improves as the number of training samples increases. Notably, for different values of d_k , the plot of the GL score against the N_k/d_k remains approximately consistent. This observation indicates that the phase transition curve for U-Net learning the MoLRG distribution depends on the ratio N_k/d_k rather than on N_k and d_k individually. When $N_k/d_k \approx 60$, GL score ≈ 1.0 suggesting that U-Net generalizes when $N_k \geq 60d_k$. This linear relationship for the phase transition differs from $N_k \geq d_k$ in Theorem 4 due to training with U-Net instead of the optimal network parameterization in Eq. (9). Nevertheless, Theorem 2 and Theorem 4 still provide valuable insights into learning distributions via diffusion models by demonstrating a similar phase transition phenomenon and confirming a linear relationship between N_k and d_k .

Learning real image data distributions with U-Net. Finally, we train diffusion models using U-Net on real image datasets AFHQ, CelebA, FFHQ, and CIFAR-10. The detailed experiment settings are deferred to Section E.3. According to [33], we define the generalization (GL) score on

real-world image dataset as follows:

$$\text{GL score} := 1 - \mathbb{P} \left(\max_{i \in [N]} [\mathcal{M}_{\text{SSCD}}(\mathbf{x}, \mathbf{y}_i)] > 0.6 \right).$$

Here, the SSCD similarity is first introduced in [74] to measure the replication between image pair $(\mathbf{x}_1, \mathbf{x}_2)$, which is defined as follows:

$$\mathcal{M}_{\text{SSCD}}(\mathbf{x}_1, \mathbf{x}_2) = \frac{\text{SSCD}(\mathbf{x}_1) \cdot \text{SSCD}(\mathbf{x}_2)}{\|\text{SSCD}(\mathbf{x}_1)\|_2 \cdot \|\text{SSCD}(\mathbf{x}_2)\|_2}$$

where $\text{SSCD}(\cdot)$ represents a neural descriptor for copy detection of images. We empirically sample 10K initial noises to estimate the probability. Intuitively, GL score measures the dissimilarity between the generated sample \mathbf{x} and all N samples \mathbf{y}_i from the training dataset $\{\mathbf{y}_i\}_{i=1}^N$. Higher GL score indicates stronger generalizability. For each data set, we train U-Net and plot the GL score against the number of training samples in Figure 5(b).

The phase transition in the real dataset is illustrated in Figure 5(b). As observed, the order in which the samples need to generalize follows the relationship: AFHQ > CelebA > FFHQ \approx CIFAR-10. Additionally, from our previous observations in Figure 3, the relationship of the intrinsic dimensions for these datasets is: AFHQ > FFHQ > CelebA \approx CIFAR-10. Both AFHQ and CelebA align well with our theoretical analysis, which indicates that more samples are required for the model to generalize as the intrinsic dimension increases.

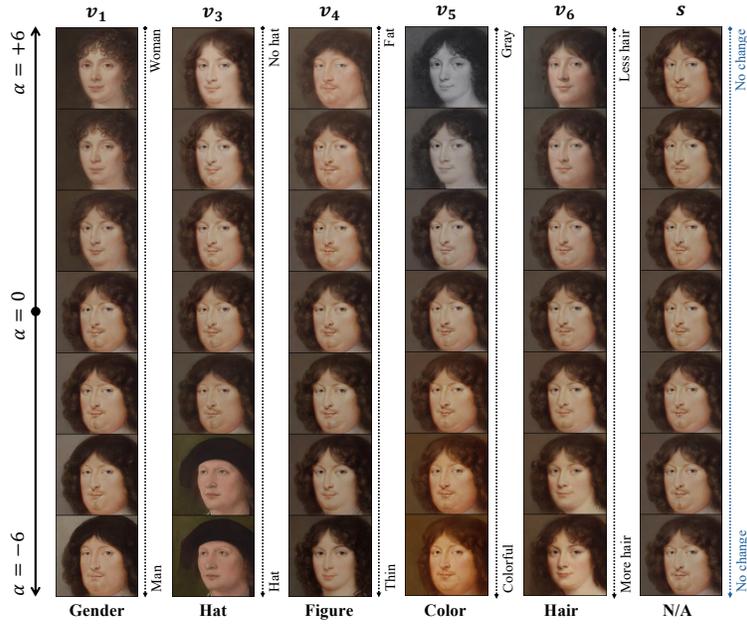
4.2 Semantic Meanings of Low-Dimensional Subspaces

In this subsection, we conduct experiments to verify the correspondence between the low-dimensional subspaces of the data distribution and the semantics of images on real datasets. We denote the Jacobian of the DAE $\mathbf{x}_\theta(\mathbf{x}_t, t)$ by $\mathbf{J}_t := \nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t) \in \mathbb{R}^{n \times n}$ and let $\mathbf{J}_t = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ be an singular value decomposition (SVD) of \mathbf{J}_t , where $r = \text{rank}(\mathbf{J}_t)$, $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_r] \in \mathcal{O}^{n \times r}$, $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_r] \in \mathcal{O}^{n \times r}$, and $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_r)$ with $\sigma_1 \geq \dots \geq \sigma_r$ being the singular values. To validate the semantic meaning of the basis vectors \mathbf{v}_i , we vary the value of α from negative to positive and visualize the resulting changes in the generated images.

In the experiments, we use a pre-trained diffusion denoising probabilistic model (DDPM) [2] on the MetFaces dataset [75]. We randomly select an image \mathbf{x}_0 from this dataset and use the reverse process of the diffusion denoising implicit model (DDIM) [76] to generate \mathbf{x}_t at $t = 0.7T$, where T denote the total number of time steps. We respectively choose the changed direction as the leading right singular vectors $\mathbf{v}_1, \mathbf{v}_3, \mathbf{v}_4, \mathbf{v}_5, \mathbf{v}_6$ and use $\tilde{\mathbf{x}}_t = \mathbf{x}_t + \alpha \mathbf{v}_i$ to generate new images with $\alpha \in [-4, 4]$ shown in Figure 6. It is observed that these singular vectors enable different semantic edits in terms of gender, hairstyle, and color of the image. For comparison, we generate a random unit vector \mathbf{s} and move \mathbf{x}_t along the direction of \mathbf{s} , where the editing strength α is the same as the semantic edits column-wise. The results are shown in the last column of Figure 6. Moving along random directions provides minimal semantic changes in the generated images, indicating that the low-dimensional subspace spanned by \mathbf{V} is non-trivial and corresponds to semantic meaningful image attributes. More experimental results can be found in Figure 9 in Section E.3.

5 Conclusion & Discussion

In this work, we studied the training loss of diffusion models to investigate when and why diffusion models can learn the underlying distribution without suffering from the curse of dimensionality.



(a)



(b)

Figure 6: **Correspondence between the singular vectors of the Jacobian of the DAE and semantic image attributes.** We use a pre-trained DDPM with U-Net on the MetFaces dataset [75]. We edit the original image \mathbf{x}_0 by changing \mathbf{x}_t into $\mathbf{x}_t + \alpha \mathbf{v}_i$, where \mathbf{v}_i is a singular vector of the Jacobian of the DAE $\mathbf{x}_\theta(\mathbf{x}_t, t)$.

Motivated by extensive empirical observations, we assumed that the underlying data distribution is a mixture of low-rank Gaussians. Specifically, we showed that minimizing the training loss is equivalent to solving the subspace clustering problem under proper network parameterization. Based on this equivalence, we further showed that the optimal solutions to the training loss can recover the underlying subspaces when the number of samples scales linearly with the intrinsic dimensionality of the data distribution. Moreover, we established the correspondence between the subspaces and semantic representations of image data. Since our studied network parameterization is not sufficiently over-parameterized, a future direction is to extend our analysis to an over-parameterized case to fully explain the transition from memorization to generalization.

Acknowledgements

P.W., H.Z., Z.Z., S.C., and Q.Q. acknowledge support from NSF CAREER CCF-2143904, NSF CCF-2212066, NSF CCF-2212326, NSF IIS 2312842, NSF IIS 2402950, a gift grant from KLA, and the MICDE Catalyst Grant. Y.M. acknowledges support from the joint Simons Foundation-NSF DMS grant #2031899, NSF IIS 2402951, and the ONR grant N00014-22-1-2102. The authors also would like to thank Laura Balzano (U. Michigan), Jeff Fessler (U. Michigan), Huikang Liu (SJTU), Dogyoon Song (UC Davis), Liyue Shen (U. Michigan), Rene Vidal (Upenn), and Zhihui Zhu (OSU) for stimulating discussions.

References

- [1] Ismail Alkhouri, Shijun Liang, Rongrong Wang, Qing Qu, and Saiprasad Ravishankar. Diffusion-based adversarial purification for robust deep mri reconstruction. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12841–12845. IEEE, 2024.
- [2] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [3] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [4] Huijie Zhang, Yifu Lu, Ismail Alkhouri, Saiprasad Ravishankar, Dogyoon Song, and Qing Qu. Improving training efficiency of diffusion models via multi-stage framework and tailored multi-decoder architectures. In *Conference on Computer Vision and Pattern Recognition 2024*, 2024. URL <https://openreview.net/forum?id=YtptmpZQ0g>.
- [5] Omer Bar-Tal, Hila Chefer, Omer Tov, Charles Herrmann, Roni Paiss, Shiran Zada, Ariel Ephrat, Junhwa Hur, Yuanzhen Li, Tomer Michaeli, et al. Lumiere: A space-time diffusion model for video generation. *arXiv preprint arXiv:2401.12945*, 2024.
- [6] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.
- [7] Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis. *Advances in Neural Information Processing Systems*, 33:17022–17033, 2020.

- [8] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. DIFFWAVE: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*, 2021.
- [9] Can Yaras, Peng Wang, Laura Balzano, and Qing Qu. Compressible dynamics in deep over-parameterized low-rank learning & adaptation. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=uDkXoZMzBv>.
- [10] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Transactions on Graphics (TOG)*, 42(1):1–13, 2022.
- [11] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22500–22510, 2023.
- [12] Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. *Advances in Neural Information Processing Systems*, 35:25683–25696, 2022.
- [13] Bowen Song, Soo Min Kwon, Zecheng Zhang, Xinyu Hu, Qing Qu, and Liyue Shen. Solving inverse problems with latent diffusion models via hard data consistency. In *The Twelfth International Conference on Learning Representations*, 2024.
- [14] Xiang Li, Soo Min Kwon, Ismail R Alkhouri, Saiprasad Ravishanka, and Qing Qu. Decoupled data consistency with diffusion purification for image restoration. *arXiv preprint arXiv:2403.06054*, 2024.
- [15] Ismail Alkhouri, Shijun Liang, Rongrong Wang, Qing Qu, and Saiprasad Ravishankar. Diffusion-based adversarial purification for robust deep mri reconstruction. *arXiv preprint arXiv:2309.05794*, 2023.
- [16] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [17] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations*, 2021.
- [18] Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- [19] Minshuo Chen, Song Mei, Jianqing Fan, and Mengdi Wang. An overview of diffusion models: Applications, guided generation, statistical rates and optimization. *arXiv preprint arXiv:2404.07771*, 2024.
- [20] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):10850–10869, 2023.

- [21] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.
- [22] Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru R Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *International Conference on Learning Representations*, 2023.
- [23] Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence for score-based generative modeling with polynomial complexity. *Advances in Neural Information Processing Systems*, 35:22870–22882, 2022.
- [24] Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. Towards faster non-asymptotic convergence for diffusion-based generative models. *arXiv preprint arXiv:2306.09251*, 2023.
- [25] Minshuo Chen, Kaixuan Huang, Tuo Zhao, and Mengdi Wang. Score approximation, estimation and distribution recovery of diffusion models on low-dimensional data. In *International Conference on Machine Learning*, pages 4672–4712. PMLR, 2023.
- [26] Kazusato Oko, Shunta Akiyama, and Taiji Suzuki. Diffusion models are minimax optimal distribution estimators. In *International Conference on Machine Learning*, pages 26517–26582. PMLR, 2023.
- [27] Xiangming Gu, Chao Du, Tianyu Pang, Chongxuan Li, Min Lin, and Ye Wang. On memorization in diffusion models. *arXiv preprint arXiv:2310.02664*, 2023.
- [28] Gowthami Somepalli, Vasu Singla, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Diffusion art or digital forgery? investigating data replication in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6048–6058, 2023.
- [29] Yuxin Wen, Yuchen Liu, Chen Chen, and Lingjuan Lyu. Detecting, explaining, and mitigating memorization in diffusion models. In *The Twelfth International Conference on Learning Representations*, 2023.
- [30] Benjamin J Zhang, Siting Liu, Wuchen Li, Markos A Katsoulakis, and Stanley J Osher. Wasserstein proximal operators describe score-based generative models and resolve memorization. *arXiv preprint arXiv:2402.06162*, 2024.
- [31] Zahra Kadkhodaie, Florentin Guth, Eero P Simoncelli, and Stéphane Mallat. Generalization in diffusion models arises from geometry-adaptive harmonic representations. In *The Twelfth International Conference on Learning Representations*, 2023.
- [32] TaeHo Yoon, Joo Young Choi, Sehyun Kwon, and Ernest K Ryu. Diffusion probabilistic models generalize when they fail to memorize. In *ICML 2023 Workshop on Structured Probabilistic Inference & Generative Modeling*, 2023.
- [33] Huijie Zhang, Jinfan Zhou, Yifu Lu, Minzhe Guo, Peng Wang, Liyue Shen, and Qing Qu. The emergence of reproducibility and consistency in diffusion models. In *Forty-first International Conference on Machine Learning*, 2023.
- [34] Andre Wibisono, Yihong Wu, and Kaylee Yingxi Yang. Optimal score estimation via empirical bayes smoothing. *arXiv preprint arXiv:2402.07747*, 2024.

- [35] Sixue Gong, Vishnu Naresh Boddeti, and Anil K Jain. On the intrinsic dimensionality of image representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3987–3996, 2019.
- [36] Phil Pope, Chen Zhu, Ahmed Abdelkader, Micah Goldblum, and Tom Goldstein. The intrinsic dimension of images and its impact on learning. In *International Conference on Learning Representations*, 2020.
- [37] Jan Pawel Stanczuk, Georgios Batzolis, Teo Deveney, and Carola-Bibiane Schönlieb. Diffusion models encode the intrinsic dimension of data manifolds. In *Forty-first International Conference on Machine Learning*, 2024.
- [38] Bradley CA Brown, Anthony L Caterini, Brendan Leigh Ross, Jesse C Cresswell, and Gabriel Loaiza-Ganem. Verifying the union of manifolds hypothesis for image data. In *The Eleventh International Conference on Learning Representations*, 2023.
- [39] Hamidreza Kamkari, Brendan Leigh Ross, Rasa Hosseinzadeh, Jesse C Cresswell, and Gabriel Loaiza-Ganem. A geometric view of data complexity: Efficient local intrinsic dimension estimation with diffusion models. *arXiv preprint arXiv:2406.03537*, 2024.
- [40] Gabriel Loaiza-Ganem, Brendan Leigh Ross, Rasa Hosseinzadeh, Anthony L Caterini, and Jesse C Cresswell. Deep generative models through the lens of the manifold hypothesis: A survey and new connections. *arXiv preprint arXiv:2404.02954*, 2024.
- [41] Arnu Pretorius, Steve Kroon, and Herman Kamper. Learning dynamics of linear denoising autoencoders. In *International Conference on Machine Learning*, pages 4141–4150. PMLR, 2018.
- [42] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [43] Pankaj K Agarwal and Nabil H Mustafa. K-means projective clustering. In *Proceedings of the 23rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pages 155–165, 2004.
- [44] René Vidal. Subspace clustering. *IEEE Signal Processing Magazine*, 28(2):52–68, 2011.
- [45] Peng Wang, Huikang Liu, Anthony Man-Cho So, and Laura Balzano. Convergence and recovery guarantees of the k-subspaces method for subspace clustering. In *International Conference on Machine Learning*, pages 22884–22918. PMLR, 2022.
- [46] Sitan Chen, Vasilis Kontonis, and Kulin Shah. Learning general gaussian mixtures with efficient score matching. *arXiv preprint arXiv:2404.18893*, 2024.
- [47] Frank Cole and Yulong Lu. Score-based generative models break the curse of dimensionality in learning a family of sub-gaussian distributions. In *The Twelfth International Conference on Learning Representations*, 2024.
- [48] Khashayar Gatmiry, Jonathan Kelner, and Holden Lee. Learning mixtures of gaussians using diffusion models. *arXiv preprint arXiv:2404.18869*, 2024.
- [49] Kulin Shah, Sitan Chen, and Adam Klivans. Learning mixtures of gaussians using the DDPM objective. *Advances in Neural Information Processing Systems*, 36:19636–19649, 2023.

- [50] Yuchen Wu, Minshuo Chen, Zihao Li, Mengdi Wang, and Yuting Wei. Theoretical insights for diffusion guidance: A case study for gaussian mixture models. In *Forty-first International Conference on Machine Learning*, 2024.
- [51] Sixu Li, Shi Chen, and Qin Li. A good score does not lead to a good generative model. *arXiv preprint arXiv:2401.04856*, 2024.
- [52] Puheng Li, Zhong Li, Huishuai Zhang, and Jiang Bian. On the generalization properties of diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024.
- [53] Mingyang Yi, Jiacheng Sun, and Zhenguo Li. On the generalization of diffusion model. *arXiv preprint arXiv:2305.14712*, 2023.
- [54] Gilad Lerman and Tyler Maunu. An overview of robust subspace recovery. *Proceedings of the IEEE*, 106(8):1380–1410, 2018.
- [55] Ehsan Elhamifar and René Vidal. Sparse subspace clustering: Algorithm, theory, and applications. *IEEE transactions on pattern analysis and machine intelligence*, 35(11):2765–2781, 2013.
- [56] Yu-Xiang Wang and Huan Xu. Noisy sparse subspace clustering. In *International Conference on Machine Learning*, pages 89–97. PMLR, 2013.
- [57] Mahdi Soltanolkotabi, Ehsan Elhamifar, Emmanuel J Candes, et al. Robust subspace clustering. *Annals of Statistics*, 42(2):669–699, 2014.
- [58] Yuanyuan Liu, LC Jiao, and Fanhua Shang. An efficient matrix factorization based low-rank representation for subspace clustering. *Pattern Recognition*, 46(1):284–292, 2013.
- [59] Haixia Liu, Jian-Feng Cai, and Yang Wang. Subspace clustering by (k, k) -sparse matrix factorization. *Inverse Problems & Imaging*, 11(3), 2017.
- [60] Gen Li and Yuantao Gu. Theory of spectral method for union of subspaces-based random geometry graph. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 6337–6345. PMLR, 2021.
- [61] Rene Vidal, Yi Ma, and Shankar Sastry. Generalized principal component analysis (GPCA). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(12):1945–1959, 2005.
- [62] 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: 26565–26577, 2022.
- [63] Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- [64] Xinlei Chen, Zhuang Liu, Saining Xie, and Kaiming He. Deconstructing denoising diffusion models for self-supervised learning. *arXiv preprint arXiv:2401.14404*, 2024.
- [65] Weilai Xiang, Hongyu Yang, Di Huang, and Yunhong Wang. Denoising diffusion autoencoders are unified self-supervised learners. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15802–15812, 2023.

- [66] 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.
- [67] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [68] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015.
- [69] Bin Xu Wang and John J Vastola. The hidden linear structure in score-based models and its application. *arXiv preprint arXiv:2311.10892*, 2023.
- [70] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015.
- [71] Vahid Kazemi and Josephine Sullivan. One millisecond face alignment with an ensemble of regression trees. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1867–1874, 2014.
- [72] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8188–8197, 2020.
- [73] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.
- [74] Ed Pizzi, Sreya Dutta Roy, Sugosh Nagavara Ravindra, Priya Goyal, and Matthijs Douze. A self-supervised descriptor for image copy detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14532–14542, 2022.
- [75] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- [76] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2020.
- [77] Per-Åke Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12:99–111, 1972.
- [78] Calvin Luo. Understanding diffusion models: A unified perspective. *arXiv preprint arXiv:2208.11970*, 2022.
- [79] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [80] Mark Rudelson and Roman Vershynin. Smallest singular value of a random rectangular matrix. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 62(12):1707–1739, 2009.

- [81] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

Appendices

In the appendix, the organization is as follows. We first provide proof details for Section 2, Section 3.1, and Section 3.2 in Section A, Section B, and Section C, respectively. Then, we present our experimental setups for Figure 3 in Section D and for Section 4 in Section E. Finally, some auxiliary results for proving the main theorems are provided in Section F.

To simplify our development, we introduce some further notation. We denote by $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ a multivariate Gaussian distribution with mean $\boldsymbol{\mu} \in \mathbb{R}^n$ and covariance $\boldsymbol{\Sigma} \succeq \mathbf{0}$. Given a Gaussian random vector $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, if $\boldsymbol{\Sigma} \succ \mathbf{0}$, with abuse of notation, we write its pdf as

$$\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) := \frac{1}{(2\pi)^{n/2} \det^{1/2}(\boldsymbol{\Sigma})} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right). \quad (22)$$

If a random vector $\mathbf{x} \in \mathbb{R}^n$ satisfies $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{U}\mathbf{U}^T)$ for some $\boldsymbol{\mu} \in \mathbb{R}^n$ and $\mathbf{U} \in \mathcal{O}^{n \times d}$, we have

$$\mathbf{x} = \boldsymbol{\mu} + \mathbf{U}\mathbf{a}, \quad (23)$$

where $\mathbf{a} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. Therefore, a mixture of low-rank Gaussians in Definition 1 can be expressed as

$$\mathbb{P}(\mathbf{x} = \mathbf{U}_k^* \mathbf{a}_k) = \pi_k, \text{ where } \mathbf{a}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{d_k}), \forall k \in [K]. \quad (24)$$

A Proofs in Section 2

A.1 Relation between Score Matching Loss and Denoiser Autoencoder Loss

To estimate $\nabla \log p_t(\mathbf{x})$, one can train a time-dependent score-based model $\mathbf{s}_\theta(\mathbf{x}, t)$ via minimizing the following objective [17]:

$$\min_{\boldsymbol{\theta}} \int_0^1 \xi_t \mathbb{E}_{\mathbf{x}_0 \sim p_{\text{data}}} \mathbb{E}_{\mathbf{x}_t | \mathbf{x}_0} \left[\|\mathbf{s}_\theta(\mathbf{x}_t, t) - \nabla \log p_t(\mathbf{x}_t | \mathbf{x}_0)\|^2 \right] dt, \quad (25)$$

where $\xi_t : [0, 1] \rightarrow \mathbb{R}^+$ is a positive weighting function. Let $\mathbf{x}_\theta(\cdot, t) : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ denote a neural network parameterized by parameters $\boldsymbol{\theta}$ to approximate $\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]$. According to the Tweedie's formula (5), $\mathbf{s}_\theta(\mathbf{x}_t, t) = (s_t \mathbf{x}_\theta(\mathbf{x}_t, t) - \mathbf{x}_t) / \gamma_t^2$ can be used to estimate score functions. Substituting this and $\nabla \log p_t(\mathbf{x}_t | \mathbf{x}_0) = (s_t \mathbf{x}_0 - \mathbf{x}_t) / \gamma_t^2$ due to (3) yields

$$\begin{aligned} & \min_{\boldsymbol{\theta}} \int_0^1 \xi_t \mathbb{E}_{\mathbf{x}_0 \sim p_{\text{data}}} \mathbb{E}_{\mathbf{x}_t | \mathbf{x}_0} \left[\left\| \frac{1}{\gamma_t^2} (s_t \mathbf{x}_\theta(\mathbf{x}_t, t) - \mathbf{x}_t) - \frac{1}{\gamma_t^2} (s_t \mathbf{x}_0 - \mathbf{x}_t) \right\|^2 \right] dt \\ & = \int_0^1 \frac{\xi_t}{s_t^2 \sigma_t^4} \mathbb{E}_{\mathbf{x}_0 \sim p_{\text{data}}} \mathbb{E}_{\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)} \left[\|\mathbf{x}_\theta(s_t \mathbf{x}_0 + \gamma_t \boldsymbol{\epsilon}, t) - \mathbf{x}_0\|^2 \right] dt, \end{aligned}$$

where the equality follows from $\mathbf{x}_t = s_t \mathbf{x}_0 + \gamma_t \boldsymbol{\epsilon}$ due to (3). Then, we obtain

$$\min_{\boldsymbol{\theta}} \int_0^1 \lambda_t \mathbb{E}_{\mathbf{x}_0 \sim p_{\text{data}}} \mathbb{E}_{\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)} \left[\|\mathbf{x}_\theta(s_t \mathbf{x}_0 + \gamma_t \boldsymbol{\epsilon}, t) - \mathbf{x}_0\|^2 \right] dt, \quad (26)$$

where $\lambda_t = \xi_t / (s_t^2 \sigma_t^4)$. However, only data points $\{\mathbf{x}^{(i)}\}_{i=1}^N$ sampled from the underlying data distribution p_{data} are available in practice. Therefore, we study the following empirical counterpart of Problem (26) over the training samples, i.e., Problem (6). We refer the reader to [31, Section 2.1] for more discussions on the denoising error of this problem.

A.2 Proof of in Lemma 1

Assuming that the underlying data distribution follows a mixture of low-rank Gaussians as defined in Definition 1, we first compute the ground-truth score function as follows.

Proposition 1. *Suppose that the underlying data distribution p_{data} follows a mixture of low-rank Gaussian distributions in Definition 1. In the forward process of diffusion models, the pdf of \mathbf{x}_t for each $t > 0$ is*

$$p_t(\mathbf{x}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \mathbf{0}, s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n), \quad (27)$$

where $\gamma_t = s_t \sigma_t$. Moreover, the score function of $p_t(\mathbf{x})$ is

$$\nabla \log p_t(\mathbf{x}) = -\frac{1}{\gamma_t^2} \left(\mathbf{x} - \frac{s_t^2}{s_t^2 + \gamma_t^2} \frac{\sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \mathbf{0}, s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}}{\sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \mathbf{0}, s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n)} \right). \quad (28)$$

Proof. Let $Y \in \{1, \dots, K\}$ be a discrete random variable that denotes the value of components of the mixture model. Note that $\gamma_t = s_t \sigma_t$. It follows from Definition 1 that $\mathbb{P}(Y = k) = \pi_k$ for each $k \in [K]$. We first compute

$$\begin{aligned} p_t(\mathbf{x}|Y = k) &= \int p_t(\mathbf{x}|Y = k, \mathbf{a}_k) \mathcal{N}(\mathbf{a}_k; \mathbf{0}, \mathbf{I}_{d_k}) d\mathbf{a}_k = \int p_t(\mathbf{x}|\mathbf{x}_0 = \mathbf{U}_k^* \mathbf{a}_k) \mathcal{N}(\mathbf{a}_k; \mathbf{0}, \mathbf{I}_{d_k}) d\mathbf{a}_k \\ &= \int \mathcal{N}(\mathbf{x}; s_t \mathbf{U}_k^* \mathbf{a}_k, \gamma_t^2 \mathbf{I}_n) \mathcal{N}(\mathbf{a}_k; \mathbf{0}, \mathbf{I}_{d_k}) d\mathbf{a}_k \\ &= \frac{1}{(2\pi)^{n/2} (2\pi)^{d_k/2} \gamma_t^n} \int \exp\left(-\frac{1}{2\gamma_t^2} \|\mathbf{x} - s_t \mathbf{U}_k^* \mathbf{a}_k\|^2\right) \exp\left(-\frac{1}{2} \|\mathbf{a}_k\|^2\right) d\mathbf{a}_k \\ &= \frac{1}{(2\pi)^{n/2} \gamma_t^n} \left(\frac{s_t^2 + \gamma_t^2}{\gamma_t^2}\right)^{-d/2} \exp\left(-\frac{1}{2\gamma_t^2} \mathbf{x}^T \left(\mathbf{I}_n - \frac{s_t^2}{s_t^2 + \gamma_t^2} \mathbf{U}_k^* \mathbf{U}_k^{*T}\right) \mathbf{x}\right) \\ &\quad \int \frac{1}{(2\pi)^{d_k/2}} \left(\frac{\gamma_t^2}{s_t^2 + \gamma_t^2}\right)^{-d/2} \exp\left(-\frac{s_t^2 + \gamma_t^2}{2\gamma_t^2} \left\| \mathbf{a}_k - \frac{s_t}{s_t^2 + \gamma_t^2} \mathbf{U}_k^{*T} \mathbf{x} \right\|^2\right) d\mathbf{a}_k \\ &= \frac{1}{(2\pi)^{n/2}} \frac{1}{\left((s_t^2 + \gamma_t^2)^d \gamma_t^{2(n-d)}\right)^{1/2}} \exp\left(-\frac{1}{2\gamma_t^2} \mathbf{x}^T \left(\mathbf{I}_n - \frac{s_t^2}{s_t^2 + \gamma_t^2} \mathbf{U}_k^* \mathbf{U}_k^{*T}\right) \mathbf{x}\right) \\ &= \frac{1}{(2\pi)^{n/2} \det^{1/2}(s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n)} \exp\left(-\frac{1}{2} \mathbf{x}^T (s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n)^{-1} \mathbf{x}\right) \\ &= \mathcal{N}(\mathbf{x}; \mathbf{0}, s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n), \end{aligned}$$

where the second equality follows from (3), the third equality uses (22), the fourth equality is due to the fact that $\langle \mathbf{x}, \mathbf{U}_k^* \mathbf{a} \rangle$ is an odd function, and the second to last equality uses $\det(s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n) = (s_t^2 + \gamma_t^2)^d \gamma_t^{2(n-d)}$ and $(s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n)^{-1} = (\mathbf{I}_n - s_t^2 / (s_t^2 + \gamma_t^2) \mathbf{U}_k^* \mathbf{U}_k^{*T}) / \gamma_t^2$ due to the matrix inversion lemma and $\mathbf{U}_k^{*T} \mathbf{U}_k^* = \mathbf{I}_{d_k}$. This, together with $\mathbb{P}(Y = k) = \pi_k$ for each $k \in [K]$, yields

$$p_t(\mathbf{x}) = \sum_{k=1}^K p_t(\mathbf{x}|Y = k) \mathbb{P}(Y = k) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \mathbf{0}, s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n).$$

Next, we directly compute

$$\begin{aligned}\nabla \log p_t(\mathbf{x}) &= \frac{\nabla p_t(\mathbf{x})}{p_t(\mathbf{x})} = \frac{\sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \mathbf{0}, s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n) \left(-\frac{1}{\gamma_t^2} \mathbf{x} + \frac{s_t^2}{\gamma_t^2 (s_t^2 + \gamma_t^2)} \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x} \right)}{\sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \mathbf{0}, s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n)} \\ &= -\frac{1}{\gamma_t^2} \left(\mathbf{x} - \frac{s_t^2}{s_t^2 + \gamma_t^2} \frac{\sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \mathbf{0}, s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}}{\sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \mathbf{0}, s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n)} \right).\end{aligned}$$

□

Proof of Lemma 1. According to (5) and Proposition 1, we compute

$$\begin{aligned}\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] &= \frac{\mathbf{x}_t + \gamma_t^2 \nabla \log p_t(\mathbf{x}_t)}{s_t} = \frac{s_t}{s_t^2 + \gamma_t^2} \frac{\sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \mathbf{0}, s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}_t}{\sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}_t; \mathbf{0}, s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n)} \\ &= \frac{s_t}{s_t^2 + \gamma_t^2} \frac{\sum_{k=1}^K \pi_k \exp\left(-\frac{1}{2\gamma_t^2} \left(\|\mathbf{x}_t\|^2 - \frac{s_t^2}{s_t^2 + \gamma_t^2} \|\mathbf{U}_k^{*T} \mathbf{x}_t\|^2\right)\right) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}_t}{\sum_{k=1}^K \pi_k \exp\left(-\frac{1}{2\gamma_t^2} \left(\|\mathbf{x}_t\|^2 - \frac{s_t^2}{s_t^2 + \gamma_t^2} \|\mathbf{U}_k^{*T} \mathbf{x}_t\|^2\right)\right)} \\ &= \frac{s_t}{s_t^2 + \gamma_t^2} \frac{\sum_{k=1}^K \pi_k \exp\left(\frac{1}{2\gamma_t^2} \frac{s_t^2}{s_t^2 + \gamma_t^2} \|\mathbf{U}_k^{*T} \mathbf{x}_t\|^2\right) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}_t}{\sum_{k=1}^K \pi_k \exp\left(\frac{1}{2\gamma_t^2} \frac{s_t^2}{s_t^2 + \gamma_t^2} \|\mathbf{U}_k^{*T} \mathbf{x}_t\|^2\right)},\end{aligned}$$

where the third equality uses (22) and $(s_t^2 \mathbf{U}_k^* \mathbf{U}_k^{*T} + \gamma_t^2 \mathbf{I}_n)^{-1} = (\mathbf{I}_n - s_t^2 / (s_t^2 + \gamma_t^2) \mathbf{U}_k^* \mathbf{U}_k^{*T}) / \gamma_t^2$ due to the matrix inversion lemma. □

B Proofs in Section 3.1

B.1 Proof of Theorem 1

Proof of Theorem 1. Plugging (11) into the integrand of (6) yields

$$\begin{aligned}\mathbb{E}_\epsilon \left[\left\| \frac{s_t}{s_t^2 + \gamma_t^2} \mathbf{U} \mathbf{U}^T \left(s_t \mathbf{x}^{(i)} + \gamma_t \boldsymbol{\epsilon} \right) - \mathbf{x}^{(i)} \right\|^2 \right] \\ &= \left\| \frac{s_t^2}{s_t^2 + \gamma_t^2} \mathbf{U} \mathbf{U}^T \mathbf{x}^{(i)} - \mathbf{x}^{(i)} \right\|^2 + \frac{(s_t \gamma_t)^2}{(s_t^2 + \gamma_t^2)^2} \mathbb{E}_\epsilon [\|\mathbf{U} \mathbf{U}^T \boldsymbol{\epsilon}\|^2] \\ &= \left\| \frac{s_t^2}{s_t^2 + \gamma_t^2} \mathbf{U} \mathbf{U}^T \mathbf{x}^{(i)} - \mathbf{x}^{(i)} \right\|^2 + \frac{(s_t \gamma_t)^2 d}{(s_t^2 + \gamma_t^2)^2},\end{aligned}$$

where the first equality follows from $\mathbb{E}_\epsilon[\langle \mathbf{x}, \boldsymbol{\epsilon} \rangle] = 0$ for any given $\mathbf{x} \in \mathbb{R}^n$ due to $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, and the second equality uses $\mathbb{E}_\epsilon[\|\mathbf{U} \mathbf{U}^T \boldsymbol{\epsilon}\|^2] = \mathbb{E}_\epsilon[\|\mathbf{U}^T \boldsymbol{\epsilon}\|^2] = \sum_{i=1}^d \mathbb{E}_\epsilon[\|\mathbf{u}_i^T \boldsymbol{\epsilon}\|^2] = d$ due to $\mathbf{U} \in \mathcal{O}^{n \times d}$ and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$. This, together with $\gamma_t = s_t \sigma_t$ and (6), yields

$$\ell(\mathbf{U}) = \frac{1}{N} \sum_{i=1}^N \int_0^1 \lambda_t \left(\|\mathbf{x}^{(i)}\|^2 - \frac{1 + 2\sigma_t^2}{(1 + \sigma_t^2)^2} \|\mathbf{U}^T \mathbf{x}^{(i)}\|^2 + \frac{\sigma_t^2 d}{(1 + \sigma_t^2)^2} \right) dt,$$

Obviously, minimizing the above function in terms of \mathbf{U} amounts to

$$\min_{\mathbf{U}^T \mathbf{U} = \mathbf{I}_d} - \int_0^1 \frac{(1 + 2\sigma_t^2)\lambda_t}{(1 + \sigma_t^2)^2} dt \frac{1}{N} \sum_{i=1}^N \|\mathbf{U}^T \mathbf{x}^{(i)}\|^2,$$

which is equivalent to Problem (12). \square

B.2 Proof of Theorem 2

Proof of Theorem 2. For ease of exposition, let

$$\mathbf{X} = [\mathbf{x}^{(1)} \ \dots \ \mathbf{x}^{(N)}] \in \mathbb{R}^{n \times N}, \ \mathbf{A} = [\mathbf{a}_1 \ \dots \ \mathbf{a}_N] \in \mathbb{R}^{d \times N}, \ \mathbf{E} = [\mathbf{e}_1 \ \dots \ \mathbf{e}_N] \in \mathbb{R}^{n \times N}.$$

Using this and (10), we obtain

$$\mathbf{X} = \mathbf{U}^* \mathbf{A} + \mathbf{E}. \quad (29)$$

Let $r_A := \text{rank}(\mathbf{A}) \leq \min\{d, N\}$ and $\mathbf{A} = \mathbf{U}_A \mathbf{\Sigma}_A \mathbf{V}_A^T$ be an singular value decomposition (SVD) of \mathbf{A} , where $\mathbf{U}_A \in \mathcal{O}^{d \times r_A}$, $\mathbf{V}_A \in \mathcal{O}^{N \times r_A}$, and $\mathbf{\Sigma}_A \in \mathbb{R}^{r_A \times r_A}$. It follows from Theorem 1 that Problem (6) with the parameterization (11) is equivalent to Problem (12).

(i) Suppose that $N \geq d$. Applying Lemma 3 with $\varepsilon = 1/(2c_1)$ to $\mathbf{A} \in \mathbb{R}^{d \times N}$, it holds with probability at least $1 - 1/2^{N-d+1} - \exp(-c_2N)$ that

$$\sigma_{\min}(\mathbf{A}) = \sigma_d(\mathbf{A}) \geq \frac{\sqrt{N} - \sqrt{d-1}}{2c_1}, \quad (30)$$

where $c_1, c_2 > 0$ are constants depending polynomially only on the Gaussian moment. This implies $r_A = d$ and $\mathbf{U}_A \in \mathcal{O}^d$. Since Problem (12) is a PCA problem, the columns of any optimal solution $\hat{\mathbf{U}} \in \mathcal{O}^{n \times d}$ consist of left singular vectors associated with the top d singular values of \mathbf{X} . This, together with Wedin's Theorem [77] and (29), yields

$$\left\| \hat{\mathbf{U}} \hat{\mathbf{U}}^T - \mathbf{U}^* \mathbf{U}^{*T} \right\|_F = \left\| \hat{\mathbf{U}} \hat{\mathbf{U}}^T - (\mathbf{U}^* \mathbf{U}_A) (\mathbf{U}^* \mathbf{U}_A)^T \right\|_F \leq \frac{2\|\mathbf{E}\|_F}{\sigma_{\min}(\mathbf{A})} = \frac{4c_1\|\mathbf{E}\|_F}{\sqrt{N} - \sqrt{d-1}}.$$

This, together with absorbing 4 into c_1 , yields (13).

(ii) Suppose that $N < d$. According to Lemma 3 with $\varepsilon = 1/(2c_1)$, it holds with probability at least $1 - 1/2^{d-N+1} - \exp(-c_2d)$ that

$$\sigma_{\min}(\mathbf{A}) = \sigma_N(\mathbf{A}) \geq \frac{\sqrt{d} - \sqrt{N-1}}{2c_1}, \quad (31)$$

where $c_1, c_2 > 0$ are constants depending polynomially only on the Gaussian moment. This implies $r_A = N$ and $\mathbf{U}_A \in \mathcal{O}^{d \times N}$. This, together with the fact that $\mathbf{A} = \mathbf{U}_A \mathbf{\Sigma}_A \mathbf{V}_A^T$ is an SVD of \mathbf{A} , yields that $\mathbf{U}^* \mathbf{A} = (\mathbf{U}^* \mathbf{U}_A) \mathbf{\Sigma}_A \mathbf{V}_A^T$ is an SVD of $\mathbf{U}^* \mathbf{A}$ with $\mathbf{U}^* \mathbf{U}_A \in \mathcal{O}^{n \times N}$. Note that $\text{rank}(\mathbf{X}) \leq N$. Let $\mathbf{X} = \mathbf{U}_X \mathbf{\Sigma}_X \mathbf{V}_X^T$ be an SVD of \mathbf{X} , where $\mathbf{U}_X \in \mathcal{O}^{n \times N}$, $\mathbf{V}_X \in \mathcal{O}^N$, and $\mathbf{\Sigma}_X \in \mathbb{R}^{N \times N}$. This, together with Wedin's Theorem [77] and (31), yields

$$\left\| \mathbf{U}_X \mathbf{U}_X^T - \mathbf{U}^* \mathbf{U}_A \mathbf{U}_A^T \mathbf{U}^{*T} \right\|_F \leq \frac{2\|\mathbf{E}\|_F}{\sigma_{\min}(\mathbf{A})} = \frac{4c_1\|\mathbf{E}\|_F}{\sqrt{d} - \sqrt{N-1}}. \quad (32)$$

Note that Problem (12) has infinite optimal solutions when $N < d$, which take the form of

$$\hat{\mathbf{U}} = [\mathbf{U}_X \ \bar{\mathbf{U}}_X] \in \mathcal{O}^{n \times d}.$$

Now, we consider that $\bar{\mathbf{U}}_X \in \mathcal{O}^{n \times (d-N)}$ is an optimal solution of the following problem:

$$\min_{\mathbf{V} \in \mathcal{O}^{n \times (d-N)}, \mathbf{U}_X^T \mathbf{V} = \mathbf{0}} \|\mathbf{V}^T \mathbf{U}^* (\mathbf{I} - \mathbf{U}_A \mathbf{U}_A^T)\|_F^2. \quad (33)$$

Then, one can verify that the rank of the following matrix is at most d :

$$\mathbf{B} := [\mathbf{U}_X \quad \mathbf{U}^* (\mathbf{I} - \mathbf{U}_A \mathbf{U}_A^T)]$$

Then, if $n \geq 2d - N$, it is easy to see that the optimal value of Problem (33) is 0. If $n < 2d - N$, the optima value is achieved at $\mathbf{V}^* = [\mathbf{V}_1^* \quad \mathbf{V}_2^*]$ with $\mathbf{V}_1^* \in \mathbb{R}^{n \times (n-d)}$ and $\mathbf{V}_2^* \in \mathbb{R}^{n \times (2d-N-n)}$ satisfying $\mathbf{V}_1^{*T} \mathbf{B} = \mathbf{0}$, which implies

$$\|\mathbf{V}^{*T} \mathbf{U}^* (\mathbf{I} - \mathbf{U}_A \mathbf{U}_A^T)\|_F^2 = \|\mathbf{V}_2^{*T} \mathbf{U}^* (\mathbf{I} - \mathbf{U}_A \mathbf{U}_A^T)\|_F^2 \leq 2d - N - n.$$

Consequently, the optimal value of Problem (33) is less than

$$\max\{0, 2d - (n + N)\} \quad (34)$$

Then, we obtain that

$$\begin{aligned} \left\| \hat{\mathbf{U}} \hat{\mathbf{U}}^T - \mathbf{U}^* \mathbf{U}^{*T} \right\|_F &= \left\| \mathbf{U}_X \mathbf{U}_X^T + \bar{\mathbf{U}}_X \bar{\mathbf{U}}_X^T - \mathbf{U}^* \mathbf{U}_A \mathbf{U}_A^T \mathbf{U}^{*T} - \mathbf{U}^* (\mathbf{I} - \mathbf{U}_A \mathbf{U}_A^T) \mathbf{U}^{*T} \right\| \\ &\geq \left\| \bar{\mathbf{U}}_X \bar{\mathbf{U}}_X^T - \mathbf{U}^* (\mathbf{I} - \mathbf{U}_A \mathbf{U}_A^T) \mathbf{U}^{*T} \right\|_F - \left\| \mathbf{U}_X \mathbf{U}_X^T - \mathbf{U}^* \mathbf{U}_A \mathbf{U}_A^T \mathbf{U}^{*T} \right\|_F \\ &\geq \sqrt{2(d-N) - 2 \max\{0, 2d - (n + N)\}} - \frac{4c_1 \|\mathbf{E}\|_F}{\sqrt{d} - \sqrt{N-1}} \\ &\geq \sqrt{2 \min\{d-N, n-d\}} - \frac{4c_1 \|\mathbf{E}\|_F}{\sqrt{d} - \sqrt{N-1}}, \end{aligned}$$

where the second inequality follows from $\bar{\mathbf{U}}_X = \mathbf{V}^*$ and (34). Then, we complete the proof. \square

C Proofs in Section 3.2

C.1 Theoretical Justification of the DAE (16)

Since $\mathbf{x}_t = s_t \mathbf{x}_0 + \gamma_t \boldsymbol{\epsilon}$, we compute

$$\mathbb{E}_{\boldsymbol{\epsilon}} [\|\mathbf{U}_k^T (s_t \mathbf{x}_0 + \gamma_t \boldsymbol{\epsilon})\|^2] = s_t^2 \|\mathbf{U}_k^T \mathbf{x}_0\|^2 + \gamma_t^2 \mathbb{E}_{\boldsymbol{\epsilon}} [\|\mathbf{U}_k^T \boldsymbol{\epsilon}\|^2] = s_t^2 \|\mathbf{U}_k^T \mathbf{x}_0\|^2 + \gamma_t^2 d,$$

where the first equality is due to $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ and $\mathbb{E}_{\boldsymbol{\epsilon}} [\langle \mathbf{U}_k^T \mathbf{x}_0, \mathbf{U}_k^T \boldsymbol{\epsilon} \rangle] = \mathbf{0}$ for each $k \in [K]$. This implies that when n is sufficiently large, we can approximate $w_k(\boldsymbol{\theta}; \mathbf{x}_t)$ in (9) well by

$$w_k(\boldsymbol{\theta}; \mathbf{x}_t) \approx \frac{\exp(\phi_t (s_t^2 \|\mathbf{U}_k^T \mathbf{x}_0\|^2 + \gamma_t^2 d))}{\sum_{l=1}^K \exp(\phi_t (s_t^2 \|\mathbf{U}_l^T \mathbf{x}_0\|^2 + \gamma_t^2 d))}.$$

This soft-max function can be further approximated by the hard-max function. Therefore, we directly obtain (17).

C.2 Proof of Theorem 3

Equipped with the above setup, we are ready to prove Theorem 3.

Proof of Theorem 3. Plugging (16) into the integrand of (6) yields

$$\begin{aligned}
& \mathbb{E}_\epsilon \left[\left\| \frac{s_t}{s_t^2 + \gamma_t^2} \sum_{k=1}^K \hat{w}_k(\boldsymbol{\theta}; \mathbf{x}^{(i)}) \mathbf{U}_k \mathbf{U}_k^T (s_t \mathbf{x}^{(i)} + \gamma_t \boldsymbol{\epsilon}) - \mathbf{x}^{(i)} \right\|^2 \right] \\
&= \left\| \frac{s_t^2}{s_t^2 + \gamma_t^2} \sum_{k=1}^K \hat{w}_k(\boldsymbol{\theta}; \mathbf{x}^{(i)}) \mathbf{U}_k \mathbf{U}_k^T \mathbf{x}^{(i)} - \mathbf{x}^{(i)} \right\|^2 + \frac{(s_t \gamma_t)^2}{(s_t^2 + \gamma_t^2)^2} \mathbb{E}_\epsilon \left[\left\| \sum_{k=1}^K \hat{w}_k(\boldsymbol{\theta}; \mathbf{x}^{(i)}) \mathbf{U}_k \mathbf{U}_k^T \boldsymbol{\epsilon} \right\|^2 \right] \\
&= \frac{s_t^2}{s_t^2 + \gamma_t^2} \sum_{k=1}^K \left(\frac{s_t^2}{s_t^2 + \gamma_t^2} \hat{w}_k^2(\boldsymbol{\theta}; \mathbf{x}^{(i)}) - 2\hat{w}_k(\boldsymbol{\theta}; \mathbf{x}^{(i)}) \right) \|\mathbf{U}_k^T \mathbf{x}^{(i)}\|^2 + \|\mathbf{x}^{(i)}\|^2 + \frac{(s_t \gamma_t)^2 d}{(s_t^2 + \gamma_t^2)^2} \sum_{k=1}^K \hat{w}_k(\boldsymbol{\theta}; \mathbf{x}^{(i)}),
\end{aligned}$$

where the first equality follows from $\mathbb{E}_\epsilon[\langle \mathbf{x}, \boldsymbol{\epsilon} \rangle] = 0$ for any fixed $\mathbf{x} \in \mathbb{R}^n$ due to $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, and the last equality uses $\mathbf{U}_k \in \mathcal{O}^{n \times d}$ and $\mathbf{U}_k^T \mathbf{U}_l = \mathbf{0}$ for all $k \neq l$. This, together with (6) and $\gamma_t = s_t \sigma_t$, yields

$$\begin{aligned}
\ell(\boldsymbol{\theta}) &= \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K \int_0^1 \frac{\lambda_t}{1 + \sigma_t^2} \left(\frac{1}{1 + \sigma_t^2} \hat{w}_k^2(\boldsymbol{\theta}; \mathbf{x}^{(i)}) - 2\hat{w}_k(\boldsymbol{\theta}; \mathbf{x}^{(i)}) \right) dt \|\mathbf{U}_k^T \mathbf{x}^{(i)}\|^2 + \\
&\quad \frac{1}{N} \int_0^1 \lambda_t dt \sum_{i=1}^N \|\mathbf{x}^{(i)}\|^2 + \left(\int_0^1 \frac{\sigma_t^2 \lambda_t}{(1 + \sigma_t^2)^2} dt \right) \frac{d}{N} \sum_{i=1}^N \sum_{k=1}^K \hat{w}_k^2(\boldsymbol{\theta}; \mathbf{x}^{(i)}).
\end{aligned}$$

According to (16), we can partition $[N]$ into $\{C_k(\boldsymbol{\theta})\}_{k=1}^K$, where $C_k(\boldsymbol{\theta})$ for each $k \in [K]$ is defined as follows:

$$C_k(\boldsymbol{\theta}) := \left\{ i \in [N] : \|\mathbf{U}_k^T \mathbf{x}^{(i)}\| \geq \|\mathbf{U}_l^T \mathbf{x}^{(i)}\|, \forall l \neq k \right\}, \forall k \in [K]. \quad (35)$$

Then, we obtain

$$\sum_{i=1}^N \sum_{k=1}^K \hat{w}_k^2(\boldsymbol{\theta}; \mathbf{x}^{(i)}) = \sum_{k=1}^K \sum_{i \in C_k(\boldsymbol{\theta})} 1 = N.$$

This, together with plugging (35) into the above loss function, yields minimizing $\ell(\boldsymbol{\theta})$ is equivalent to minimizing

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K \int_0^1 \frac{\lambda_t}{1 + \sigma_t^2} \left(\frac{1}{1 + \sigma_t^2} \hat{w}_k^2(\boldsymbol{\theta}; \mathbf{x}^{(i)}) - 2\hat{w}_k(\boldsymbol{\theta}; \mathbf{x}^{(i)}) \right) dt \|\mathbf{U}_k^T \mathbf{x}^{(i)}\|^2 \\
&= \left(\int_0^1 \frac{\lambda_t}{1 + \sigma_t^2} \left(\frac{1}{1 + \sigma_t^2} - 2 \right) dt \right) \frac{1}{N} \sum_{k=1}^K \sum_{i \in C_k(\boldsymbol{\theta})} \|\mathbf{U}_k^T \mathbf{x}^{(i)}\|^2.
\end{aligned}$$

Since $\frac{\lambda_t}{1 + \sigma_t^2} \left(\frac{1}{1 + \sigma_t^2} - 2 \right) < 0$ for all $t \in [0, 1]$, minimizing the above function is equivalent to

$$\max_{\boldsymbol{\theta}} \frac{1}{N} \sum_{k=1}^K \sum_{i \in C_k(\boldsymbol{\theta})} \|\mathbf{U}_k^T \mathbf{x}^{(i)}\|^2 \quad \text{s.t. } [\mathbf{U}_1 \dots \mathbf{U}_K] \in \mathcal{O}^{n \times dK}.$$

Then, we complete the proof. \square

C.3 Proof of Theorem 4

Proof of Theorem 4. For ease of exposition, let $\delta := \max\{\|\mathbf{e}_i\| : i \in [N]\}$,

$$f(\boldsymbol{\theta}) := \sum_{k=1}^K \sum_{i \in C_k(\boldsymbol{\theta})} \|\mathbf{U}_k^T \mathbf{x}^{(i)}\|^2,$$

and for each $k \in [K]$,

$$C_k^* := \left\{ i \in [N] : \mathbf{x}^{(i)} = \mathbf{U}_k^* \mathbf{a}_i + \mathbf{e}_i \right\}.$$

Suppose that (50) and (51) hold with $\mathbf{V} = \hat{\mathbf{U}}_k$ for all $i \in [N]$ and $k \neq l \in [K]$, which happens with probability $1 - 2K^2N^{-1}$ according to Lemma 5. This implies that for all $i \in [N]$ and $k \neq l \in [K]$,

$$\sqrt{d} - (2\sqrt{\log N} + 2) \leq \|\mathbf{a}_i\| \leq \sqrt{d} + (2\sqrt{\log N} + 2), \quad (36)$$

$$\|\hat{\mathbf{U}}_k^T \mathbf{U}_l^*\|_F - (2\sqrt{\log N} + 2) \leq \|\hat{\mathbf{U}}_k^T \mathbf{U}_l^* \mathbf{a}_i\| \leq \|\hat{\mathbf{U}}_k^T \mathbf{U}_l^*\|_F + (2\sqrt{\log N} + 2). \quad (37)$$

Recall that the underlying basis matrices are denoted by $\boldsymbol{\theta}^* = \{\mathbf{U}_k^*\}_{k=1}^K$ and the optimal basis matrices are denoted by $\hat{\boldsymbol{\theta}} = \{\hat{\mathbf{U}}_k\}_{k=1}^K$.

First, we claim that $C_k(\boldsymbol{\theta}^*) = C_k^*$ for each $k \in [K]$. Indeed, for each $i \in C_k^*$, we compute

$$\|\mathbf{U}_k^{*T} \mathbf{x}^{(i)}\| = \|\mathbf{U}_k^{*T} (\mathbf{U}_k^* \mathbf{a}_i + \mathbf{e}_i)\| = \|\mathbf{a}_i + \mathbf{U}_k^{*T} \mathbf{e}_i\| \geq \|\mathbf{a}_i\| - \|\mathbf{e}_i\|, \quad (38)$$

$$\|\mathbf{U}_l^{*T} \mathbf{x}^{(i)}\| = \|\mathbf{U}_l^{*T} (\mathbf{U}_k^* \mathbf{a}_i + \mathbf{e}_i)\| = \|\mathbf{U}_l^{*T} \mathbf{e}_i\| \leq \|\mathbf{e}_i\|, \quad \forall l \neq k. \quad (39)$$

This, together with (36) and $\|\mathbf{e}_i\| < (\sqrt{d} - 2\sqrt{\log N})/2$, implies $\|\mathbf{U}_k^{*T} \mathbf{x}^{(i)}\| \geq \|\mathbf{U}_l^{*T} \mathbf{x}^{(i)}\|$ for all $l \neq k$. Therefore, we have $i \in C_k(\boldsymbol{\theta}^*)$ due to (35). Therefore, we have $C_k^* \subseteq C_k(\boldsymbol{\theta}^*)$ for each $k \in [K]$. This, together with the fact that they respectively denote a partition of $[N]$, yields $C_k(\boldsymbol{\theta}^*) = C_k^*$ for each $k \in [K]$. Now, we compute

$$\begin{aligned} f(\boldsymbol{\theta}^*) &= \sum_{k=1}^K \sum_{i \in C_k^*} \|\mathbf{U}_k^{*T} \mathbf{x}^{(i)}\|^2 = \sum_{k=1}^K \sum_{i \in C_k^*} \|\mathbf{a}_i + \mathbf{U}_k^{*T} \mathbf{e}_i\|^2 \\ &= \sum_{i=1}^N \|\mathbf{a}_i\|^2 + 2 \sum_{k=1}^K \sum_{i \in C_k^*} \langle \mathbf{a}_i, \mathbf{U}_k^{*T} \mathbf{e}_i \rangle + \sum_{k=1}^K \sum_{i \in C_k^*} \|\mathbf{U}_k^{*T} \mathbf{e}_i\|^2. \end{aligned} \quad (40)$$

Next, we compute

$$\begin{aligned} f(\hat{\boldsymbol{\theta}}) &= \sum_{k=1}^K \sum_{i \in C_k(\hat{\boldsymbol{\theta}})} \|\hat{\mathbf{U}}_k^T \mathbf{x}^{(i)}\|^2 = \sum_{l=1}^K \sum_{k=1}^K \sum_{i \in C_k(\hat{\boldsymbol{\theta}}) \cap C_l^*} \|\hat{\mathbf{U}}_k^T (\mathbf{U}_l^* \mathbf{a}_i + \mathbf{e}_i)\|^2 \\ &= \sum_{l=1}^K \sum_{k=1}^K \sum_{i \in C_k(\hat{\boldsymbol{\theta}}) \cap C_l^*} \left(\|\hat{\mathbf{U}}_k^T \mathbf{U}_l^* \mathbf{a}_i\|^2 + 2\langle \mathbf{a}_i, \mathbf{U}_l^{*T} \hat{\mathbf{U}}_k^T \mathbf{e}_i \rangle \right) + \sum_{k=1}^K \sum_{i \in C_k(\hat{\boldsymbol{\theta}})} \|\hat{\mathbf{U}}_k^T \mathbf{e}_i\|^2. \end{aligned}$$

This, together with $f(\hat{\boldsymbol{\theta}}) \geq f(\boldsymbol{\theta}^*)$ and (40), yields

$$\begin{aligned}
\sum_{i=1}^N \|\mathbf{a}_i\|^2 - \sum_{l=1}^K \sum_{k=1}^K \sum_{i \in C_k(\hat{\boldsymbol{\theta}}) \cap C_l^*} \|\hat{\mathbf{U}}_k^T \mathbf{U}_l^* \mathbf{a}_i\|^2 &\leq \sum_{l=1}^K \sum_{k=1}^K \sum_{i \in C_k(\hat{\boldsymbol{\theta}}) \cap C_l^*} 2\langle \mathbf{a}_i, \mathbf{U}_l^{*T} \hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^T \mathbf{e}_i \rangle + \\
&\sum_{k=1}^K \sum_{i \in C_k(\hat{\boldsymbol{\theta}})} \|\hat{\mathbf{U}}_k^T \mathbf{e}_i\|^2 - 2 \sum_{k=1}^K \sum_{i \in C_k^*} \langle \mathbf{a}_i, \mathbf{U}_k^{*T} \mathbf{e}_i \rangle - \sum_{k=1}^K \sum_{i \in C_k^*} \|\mathbf{U}_k^{*T} \mathbf{e}_i\|^2 \\
&\leq 4\delta \sum_{i=1}^N \|\mathbf{a}_i\| + N\delta^2 \leq 6\delta N\sqrt{d} + N\delta^2,
\end{aligned} \tag{41}$$

where the second inequality follows from $\|\mathbf{e}_i\| \leq \delta$ for all $i \in [N]$ and $\mathbf{U}_k^*, \hat{\mathbf{U}}_k \in \mathcal{O}^{n \times d}$ for all $k \in [K]$, and the last inequality uses (36).

For ease of exposition, let $N_{kl} := |C_k(\hat{\boldsymbol{\theta}}) \cap C_l^*|$. According to the pigeonhole principle, there exists a permutation $\pi : [K] \rightarrow [K]$ such that there exists $k \in [K]$ such that $N_{\pi(k)k} \geq N/K^2$. This, together with (41), yields

$$\begin{aligned}
6\delta N\sqrt{d} + N\delta^2 &\geq \sum_{i \in C_{\pi(k)}(\hat{\boldsymbol{\theta}}) \cap C_k^*} \left(\|\mathbf{a}_i\|^2 - \|\hat{\mathbf{U}}_{\pi(k)}^T \mathbf{U}_k^* \mathbf{a}_i\|^2 \right) \\
&= \langle \mathbf{I} - \mathbf{U}_k^{*T} \hat{\mathbf{U}}_{\pi(k)} \hat{\mathbf{U}}_{\pi(k)}^T \mathbf{U}_k^*, \sum_{i \in C_{\pi(k)}(\hat{\boldsymbol{\theta}}) \cap C_k^*} \mathbf{a}_i \mathbf{a}_i^T \rangle.
\end{aligned} \tag{42}$$

According to Lemma 6 and $N_{\pi(k)k} \geq N/K^2$, it holds with probability at least $1 - 2K^4N^{-2}$ that

$$\left\| \frac{1}{N_{\pi(k)k}} \sum_{i \in C_{\pi(k)}(\hat{\boldsymbol{\theta}}) \cap C_k^*} \mathbf{a}_i \mathbf{a}_i^T - \mathbf{I} \right\| \leq \frac{9(\sqrt{d} + \sqrt{\log(N_{\pi(k)k})})}{\sqrt{N_{\pi(k)k}}}.$$

This, together with the Weyl's inequality, yields

$$\begin{aligned}
\lambda_{\min} \left(\sum_{i \in C_{\pi(k)}(\hat{\boldsymbol{\theta}}) \cap C_k^*} \mathbf{a}_i \mathbf{a}_i^T \right) &\geq N_{\pi(k)k} - 9\sqrt{N_{\pi(k)k}} \left(\sqrt{d} + \sqrt{\log(N_{\pi(k)k})} \right) \\
&\geq \frac{N}{K^2} - \frac{9\sqrt{N}}{K} \left(\sqrt{d} + \sqrt{\log N} \right) \geq \frac{N}{2K^2},
\end{aligned}$$

where the second inequality follows from $N/K^2 \leq N_{\pi(k)k} \leq N$ and the last inequality is due to $\sqrt{N} \geq 18K(\sqrt{d} + \sqrt{\log N})$. Using this and Lemma 7, we obtain

$$\begin{aligned}
&\langle \mathbf{I} - \mathbf{U}_k^{*T} \hat{\mathbf{U}}_{\pi(k)} \hat{\mathbf{U}}_{\pi(k)}^T \mathbf{U}_k^*, \sum_{i \in C_{\pi(k)}(\hat{\boldsymbol{\theta}}) \cap C_k^*} \mathbf{a}_i \mathbf{a}_i^T \rangle \\
&\geq \lambda_{\min} \left(\sum_{i \in C_{\pi(k)}(\hat{\boldsymbol{\theta}}) \cap C_k^*} \mathbf{a}_i \mathbf{a}_i^T \right) \text{Tr} \left(\mathbf{I} - \mathbf{U}_k^{*T} \hat{\mathbf{U}}_{\pi(k)} \hat{\mathbf{U}}_{\pi(k)}^T \mathbf{U}_k^* \right) \\
&\geq \frac{N}{2K^2} \text{Tr} \left(\mathbf{I} - \mathbf{U}_k^{*T} \hat{\mathbf{U}}_{\pi(k)} \hat{\mathbf{U}}_{\pi(k)}^T \mathbf{U}_k^* \right).
\end{aligned}$$

This, together with (42), implies

$$\mathrm{Tr} \left(\mathbf{I} - \mathbf{U}_k^{\star T} \hat{\mathbf{U}}_{\pi(k)} \hat{\mathbf{U}}_{\pi(k)}^T \mathbf{U}_k^{\star} \right) \leq 2K^2 \left(6\delta\sqrt{d} + \delta^2 \right).$$

Using this and $[\mathbf{U}_1^{\star}, \dots, \mathbf{U}_K^{\star}] \in \mathcal{O}^{n \times dK}$, we obtain

$$\begin{aligned} \sum_{l \neq k} \|\hat{\mathbf{U}}_{\pi(k)}^T \mathbf{U}_l^{\star}\|_F^2 &= \mathrm{Tr} \left(\sum_{l \neq k} \hat{\mathbf{U}}_{\pi(k)}^T \mathbf{U}_l^{\star} \mathbf{U}_l^{\star T} \hat{\mathbf{U}}_{\pi(k)} \right) \leq \mathrm{Tr} \left(\mathbf{I} - \hat{\mathbf{U}}_{\pi(k)}^T \mathbf{U}_k^{\star} \mathbf{U}_k^{\star T} \hat{\mathbf{U}}_{\pi(k)} \right) \\ &\leq 2K^2 \left(6\delta\sqrt{d} + \delta^2 \right) \leq \frac{3d}{4}, \end{aligned} \quad (43)$$

where the last inequality follows $\delta \leq \sqrt{d}/(24K^2)$. According to (41), we have

$$\begin{aligned} 6\delta N\sqrt{d} + N\delta^2 &\geq \sum_{l \neq k} \sum_{i \in C_{\pi(k)}(\hat{\boldsymbol{\theta}}) \cap C_l^{\star}} \left(\|\mathbf{a}_i\|^2 - \|\hat{\mathbf{U}}_{\pi(k)}^T \mathbf{U}_l^{\star} \mathbf{a}_i\|^2 \right) \\ &\geq \sum_{l \neq k} N_{\pi(k)l} \left((\sqrt{d} - \alpha)^2 - \left(\|\hat{\mathbf{U}}_{\pi(k)}^T \mathbf{U}_l^{\star}\|_F + \alpha \right)^2 \right) \geq \frac{d}{8} \sum_{l \neq k} N_{\pi(k)l}, \end{aligned}$$

where the second inequality uses (36) and (37), and the last inequality follows from $d \gtrsim \log N$. Therefore, we have for each $k \in [K]$,

$$\sum_{l \neq k} N_{\pi(k)l} \leq \frac{48\delta N\sqrt{d} + 8\delta^2 N}{d} < 1,$$

where the last inequality uses $\delta \lesssim \sqrt{d/N}$. This implies $N_{\pi(l)k} = 0$ for all $l \neq k$, and thus $C_{\pi(k)}(\hat{\boldsymbol{\theta}}) \subseteq C_k^{\star}$. Using the same argument, we can show that $C_{\pi(l)}(\hat{\boldsymbol{\theta}}) \subseteq C_l^{\star}$ for each $l \neq k$. Therefore, we have $C_{\pi(k)}(\hat{\boldsymbol{\theta}}) = C_k^{\star}$ for each $k \in [K]$. In particular, using the union bound yields event holds with probability at least $1 - 2K^2N^{-1}$. Based on the above optimal assignment, we can further show:

(i) Suppose that $N_k \geq d$ for each $k \in [K]$. This, together with (i) in Theorem 2 and $N_k \geq d$, yields (19).

(ii) Suppose that there exists $k \in [K]$ such that $N_k < d$. This, together with (ii) in Theorem 2 and $N_k \geq d$, yields (20).

Finally, applying the union bound yields the probability of these events. \square

D Experimental Setups in Section 2.2

In this section, we provide detailed setups for the experiments in Section 2.2. These experiments aim to validate the assumptions that real-world image data satisfies a mixture of low-rank Gaussians and that the DAE is parameterized according to (9). To begin, we show that $\nabla_{\mathbf{x}_i} \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_i]$ is of low rank when p_{data} follows a mixture of low-rank Gaussians and $\sum_{k=1}^K d_k \leq n$, where n is the ambient dimension of training samples.

D.1 Verification of Mixture of Low-Rank Gaussian Data Distribution

In this subsection, we demonstrate that a mixture of low-rank Gaussians is a reasonable and insightful model for approximating real-world image data distribution. To begin, we show that $\nabla_{\mathbf{x}_t} \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]$ is of low rank when p_{data} follows a mixture of low-rank Gaussians with $\sum_{k=1}^K d_k \leq n$, where n is the dimension of training samples.

Lemma 2. *Suppose that the data distribution p_{data} follows a mixture of low-rank Gaussian distributions as defined in Definition 1. For all $t \in [0, 1]$, it holds that*

$$\min_{k \in [K]} d_k \leq \text{rank}(\nabla_{\mathbf{x}_t} \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]) \leq \sum_{k=1}^K d_k. \quad (44)$$

Proof. For ease of exposition, let

$$h_k(\mathbf{x}_t) := \exp(\phi_t \|\mathbf{U}_k^{*T} \mathbf{x}_t\|^2), \quad \forall k \in [K].$$

Obviously, we have

$$\nabla h_k(\mathbf{x}_t) := 2\phi_t \exp(\phi_t \|\mathbf{U}_k^{*T} \mathbf{x}_t\|^2) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}_t = 2\phi_t h_k(\mathbf{x}_t) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}_t. \quad (45)$$

According to Lemma 1, we have

$$\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = \frac{s_t}{s_t^2 + \gamma_t^2} f(\mathbf{x}_t), \quad \text{where } f(\mathbf{x}_t) := \frac{\sum_{k=1}^K \pi_k h_k(\mathbf{x}_t) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}_t}{\sum_{k=1}^K \pi_k h_k(\mathbf{x}_t)}.$$

Then, we compute

$$\begin{aligned} \nabla_{\mathbf{x}_t} f(\mathbf{x}_t) &= \frac{1}{\sum_{k=1}^K \pi_k h_k(\mathbf{x}_t)} \left(2\phi_t \sum_{k=1}^K \pi_k h_k(\mathbf{x}_t) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}_t \mathbf{x}_t^T \mathbf{U}_k^* \mathbf{U}_k^{*T} + \sum_{k=1}^K \pi_k h_k(\mathbf{x}_t) \mathbf{U}_k^* \mathbf{U}_k^{*T} \right) \\ &\quad - \frac{2\phi_t}{\left(\sum_{k=1}^K \pi_k h_k(\mathbf{x}_t) \right)^2} \left(\sum_{k=1}^K \pi_k h_k(\mathbf{x}_t) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}_t \right) \left(\sum_{k=1}^K \pi_k h_k(\mathbf{x}_t) \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}_t \right)^T \\ &= \frac{1}{\sum_{k=1}^K \pi_k h_k(\mathbf{x}_t)} \sum_{k=1}^K \pi_k h_k(\mathbf{x}_t) (2\phi_t \mathbf{U}_k^* \mathbf{U}_k^{*T} \mathbf{x}_t \mathbf{x}_t^T + \mathbf{I}) \mathbf{U}_k^* \mathbf{U}_k^{*T} - \\ &\quad \frac{2\phi_t}{\left(\sum_{k=1}^K \pi_k h_k(\mathbf{x}_t) \right)^2} \left(\sum_{k=1}^K \pi_k h_k(\mathbf{x}_t) \mathbf{U}_k^* \mathbf{U}_k^{*T} \right) \mathbf{x}_t \mathbf{x}_t^T \left(\sum_{k=1}^K \pi_k h_k(\mathbf{x}_t) \mathbf{U}_k^* \mathbf{U}_k^{*T} \right). \end{aligned}$$

This directly yields (44) for all $t \in [0, 1]$. \square

Now, we conduct experiments to illustrate that diffusion models trained on real-world image datasets exhibit similar low-rank properties to those described in the above proposition. Provided that the DAE $\mathbf{x}_\theta(\mathbf{x}_t, t)$ is applied to estimate $\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]$, we estimate the rank of the Jacobian of the DAE, i.e., $\nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t)$, on the real-world data distribution, where θ denotes the parameters of U-Net architecture trained on the real dataset. Also, this estimation is based on the findings in [78, 33] that under the training loss in Equation (6), the DAE $\mathbf{x}_\theta(\mathbf{x}_t, t)$ converge to $\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]$ as the number of training samples increases on the real data. We evaluate the numerical rank of the

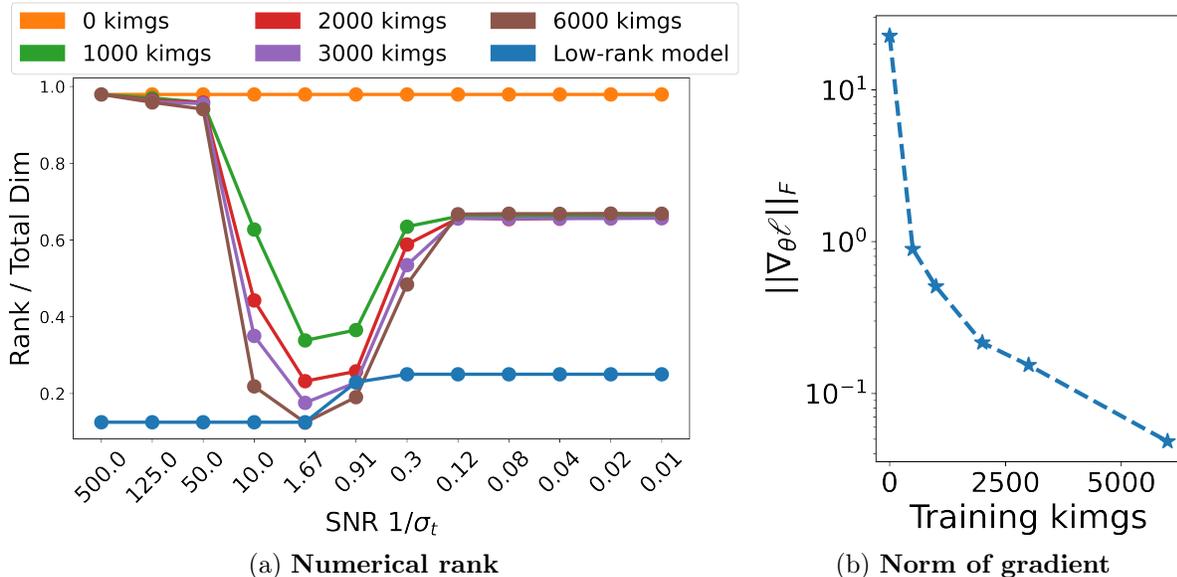


Figure 7: (a) **Numerical rank of $\nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t)$ at all time of diffusion models.** Problem (6) is trained with the DAE $\mathbf{x}_\theta(\cdot, t)$ parameterized according to (9) and U-Net on the training samples generated by the mixture of low-rank Gaussian distribution. The x -axis is the SNR and the y -axis is the numerical rank of $\nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t)$ over the ambient dimension n , i.e., $\text{rank}(\nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t))/n$. Here, *kimgs* denotes the number of samples used for training, which equals to training iterations times batch size of training samples. (b) **Convergence of gradient norm of the training loss:** The x -axis is kimgs (see Eq. (47)), and the y -axis is the gradient norm of the training loss.

Jacobian of the DAE on four different datasets: CIFAR-10 [67], CelebA [70], FFHQ [71] and AFHQ [72], where the ambient dimension $n = 3072$ for all datasets.

Given a random initial noise $\mathbf{x}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, diffusion models generate a sequence of images $\{\mathbf{x}_t\}$ according to the reverse SDE in Eq. (4). Along the sampling trajectory $\{\mathbf{x}_t\}$, we calculate the Jacobian $\nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t)$ and compute its numerical rank via

$$\text{rank}(\nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t)) := \arg \min \left\{ r \in [1, n] : \frac{\sum_{i=1}^r \sigma_i^2(\nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t))}{\sum_{i=1}^n \sigma_i^2(\nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t))} > \eta^2 \right\}. \quad (46)$$

In our experiments, we set $\eta = 0.99$. In the implementation, we utilize the Elucidating Diffusion Model (EDM) with the EDM noise scheduler [62] and DDPM++ architecture [76]. Moreover, we employ an 18-step Heun’s solver for sampling and present the results for 12 of these steps. For each dataset, we random sample 15 initial noise \mathbf{x}_1 , calculate the mean of $\text{rank}(\nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t))$ along the trajectory $\{\mathbf{x}_t\}$, and plot ratio of the numerical rank over the ambient dimension against the signal-noise-ratio (SNR) $1/\sigma_t$ in Figure 3, where σ_t is defined in Eq. (3).

D.2 Verification of Low-Rank Network Parameterization

In this subsection, we empirically investigate the properties of U-Net architectures in diffusion models and validate the simplification of the network architecture to Eq. (9). Based on the results in Section D.1, we use a mixture of low-rank Gaussian distributions for experiments. Here, we set $K = 2$, $n = 48$, $d_1 = d_2 = 6$, $\pi_1 = \pi_2 = 0.5$, and $N = 1000$ for the data model Definition 1. Moreover, We use the EDM noise scheduler and 18-step Heun’s solver for both the U-Net and our proposed parameterization (9). To adapt the structure of the U-Net, we reshape each training

Algorithm 1 SGD for optimizing the training loss (6)

Input: Training samples $\{\mathbf{x}^{(i)}\}_{i=1}^N$

for $j = 0, 1, 2, \dots, J$ **do**

 Randomly select $\{(i_m, t_m)\}_{m=1}^M$, where $i_m \in [N]$ and $t_m \in (0, 1)$ and a noise $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

 Take a gradient step

$$\boldsymbol{\theta}^{j+1} \leftarrow \boldsymbol{\theta}^j - \frac{\eta}{M} \sum_{m \in [M]} \nabla_{\boldsymbol{\theta}} \left\| \mathbf{x}_{\boldsymbol{\theta}^j}(s_{t_m} \mathbf{x}^{(i_m)} + \gamma_{t_m} \boldsymbol{\epsilon}, t_m) - \mathbf{x}^{(i_m)} \right\|^2$$

end for

sample into a 3D tensor with dimensions $4 \times 4 \times 3$, treating it as an image. Here, we use DDPM++ based diffusion models with a U-Net architecture. In each iteration, we randomly sampled a batch of image $\{\mathbf{x}^{(j)}\}_{j=1}^{\text{bs}} \subseteq \{\mathbf{x}^{(i)}\}_{i=1}^N$, along with a timestep $t^{(j)}$ and a noise $\boldsymbol{\epsilon}^{(j)}$ for each image in the batch to optimize the training loss $\ell(\boldsymbol{\theta})$. We define

$$\text{kings} = \text{bs} \times \frac{\text{training iterations}}{1000} \quad (47)$$

to represent the total samples used for training. Here, we pick up the specific model trained under 500 kings, 1000 kings, 2000 kings, and 6000 kings for evaluation, as shown in Figure 7(a).

We plot the numerical ranks of $\nabla_{\mathbf{x}_t} \mathbf{x}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)$ for both our proposed parameterization in (9) and for the U-Net architecture in Figure 3(b). According to Lemma 2, it holds that $6 \leq \text{rank}(\nabla_{\mathbf{x}_t} \mathbf{x}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)) \leq 12$. This corresponds to the blue curve in Figure 3(b). To supplement our result in Figure 3(b), we further plot the numerical rank against SNR at different training iterations in Figure 7(a) and gradient norm of the objective against training iterations in Figure 7(b). We observe that with the training kings increases, the gradient for the U-Net $\|\nabla_{\boldsymbol{\theta}} \ell\|_F$ decrease smaller than 10^{-1} and the rank ratio of $\nabla_{\mathbf{x}_t} \mathbf{x}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)$ trained from U-Net gradually be close to the rank ratio from the low-rank model in the middle of the SNR ([0.91, 10.0]).

E Experimental Setups in Section 4

We use a CPU to optimize Problem (6) for the setting in Section E.1. For the settings in Section E.2 and Section E.3, we employ a single A40 GPU with 48 GB memory to optimize Problem (6).

E.1 Learning the MoLRG distribution with the theoretical parameterization

Here, we present the stochastic gradient descent (SGD) algorithm for solving Problem (6) as follows:

Now, we specify how to choose the parameters of the SGD in our implementation. We divide the time interval $[0, 1]$ into 64 time steps. When $K = 1$, we set the learning rate $\eta = 10^{-4}$, batch size $M = 128N_k$, and number of iterations $J = 10^4$. When $K = 2$, we set the learning rate $\eta = 2 \times 10^{-5}$, batch size $M = 1024$, number of iterations $J = 10^5$. In particular, when $K = 2$, we use the following tailor-designed initialization $\boldsymbol{\theta}^0 = \{\mathbf{U}_k^0\}$ to improve the convergence of the SGD:

$$\mathbf{U}_k^0 = \mathbf{U}_k^* + 0.2\boldsymbol{\Delta}, \quad k \in \{1, 2\}, \quad (48)$$

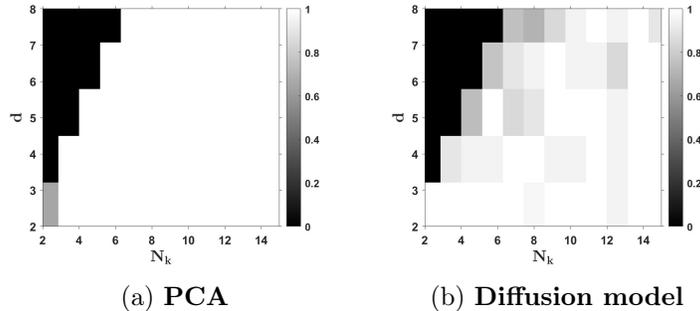


Figure 8: **Phase transition of learning the MoLRG distribution when $K = 3$.** The x -axis is the number of training samples and y -axis is the dimension of subspaces. We apply a subspace clustering method and train diffusion models for solving Problems (18) and (6), visualizing the results in (a) and (b), respectively.

where $\Delta \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$. We calculate the success rate as follows. If the returned subspace basis matrices $\{\mathbf{U}_k\}_{k=1}^K$ satisfy

$$\frac{1}{K} \sum_{k=1}^K \|\mathbf{U}_{\Pi(k)} \mathbf{U}_{\Pi(k)}^T - \mathbf{U}_k^* \mathbf{U}_k^{*T}\| \leq 0.5$$

for some permutation $\Pi : [K] \rightarrow [K]$, it is considered successful.

E.2 Learning the MoLRG distribution with U-Net

In our implementation, we set the total dimension of MoLRG as $n = 48$ and the number of training samples $N_{\text{eval}} = 1000$. To train the U-Net, we use the stochastic gradient descent in Algorithm 1. We use DDPM++ architecture [17] for the U-Net and EDM [62] noise scheduler. We set the learning rate 10^{-3} , batch size 64, and number of iterations $J = 10^4$.

E.3 Learning real-world image data distributions with U-Net

To train diffusion models for real-world image datasets, we use the DDPM++ architecture [17] for the U-Net and variance preserving (VP) [17] noise scheduler. The U-Net is trained using the Adam optimizer [79], a variant of SGD in Algorithm 1. We set the learning rate $\eta = 10^{-3}$, batch size $M = 512$, and the total number of iterations 10^5 .

E.4 Correspondence between low-dimensional subspaces and image semantics

We denote the Jacobian of the DAE $\mathbf{x}_\theta(\mathbf{x}_t, t)$ by $\mathbf{J}_t := \nabla_{\mathbf{x}_t} \mathbf{x}_\theta(\mathbf{x}_t, t) \in \mathbb{R}^{n \times n}$ and let $\mathbf{J}_t = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$ be an singular value decomposition (SVD) of \mathbf{J}_t , where $r = \text{rank}(\mathbf{J}_t)$, $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_r] \in \mathcal{O}^{n \times r}$, $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_r] \in \mathcal{O}^{n \times r}$, and $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_r)$ with $\sigma_1 \geq \dots \geq \sigma_r$ being the singular values. According to the results in Figure 3, it is observed that \mathbf{J}_t is low rank, i.e., $r \ll n$. Now, we compute the first-order approximation of $\mathbf{x}_\theta(\mathbf{x}_t, t)$ along the direction of $\mathbf{v}_i \in \mathbb{R}^n$, where \mathbf{v}_i is the i -th right singular vector of \mathbf{J}_t :

$$\mathbf{x}_\theta(\mathbf{x}_t + \alpha \mathbf{v}_i, t) \approx \mathbf{x}_\theta(\mathbf{x}_t, t) + \alpha \mathbf{J}_t \mathbf{v}_i = \mathbf{x}_\theta(\mathbf{x}_t, t) + \alpha \sigma_i \mathbf{u}_i,$$

where the last equality follows from $\mathbf{J}_t \mathbf{v}_i = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T \mathbf{v}_i = \alpha \sigma_i \mathbf{u}_i$. To validate the semantic meaning of the basis \mathbf{v}_i , we vary the value of α from negative to positive and visualize the resulting changes in the generated images. Figures 6 and 9(a, c) illustrate some real examples.

In the experiments, we use a pre-trained diffusion denoising probabilistic model (DDPM) [2] on the MetFaces dataset [75]. We randomly select an image \mathbf{x}_0 from this dataset and use the reverse process of the diffusion denoising implicit model (DDIM) [76] to generate \mathbf{x}_t at $t = 0.7T$ (ablation studies for $t = 0.1T$ and $0.9T$ are shown in Figure 9(b)), where T denote the total number of time steps. We respectively choose the changed direction as the leading right singular vectors $\mathbf{v}_1, \mathbf{v}_3, \mathbf{v}_4, \mathbf{v}_5, \mathbf{v}_6$ and use $\tilde{\mathbf{x}}_t = \mathbf{x}_t + \alpha \mathbf{v}_i$ to generate new images with $\alpha \in [-6, 6]$ shown in Figures 6 and 9(a, c).

F Auxiliary Results

First, we present a probabilistic result to prove Theorem 2, which provides an optimal estimate of the small singular values of a matrix with i.i.d. Gaussian entries. This lemma is proved in [80, Theorem 1.1].

Lemma 3. *Let \mathbf{A} be an $m \times n$ random matrix, where $m \geq n$, whose elements are independent copies of a subgaussian random variable with mean zero and unit variance. It holds for every $\varepsilon > 0$ that*

$$\mathbb{P}(\sigma_{\min}(\mathbf{A}) \geq \varepsilon(\sqrt{m} - \sqrt{n-1})) \geq 1 - (c_1\varepsilon)^{m-n+1} - \exp(-c_2m),$$

where $c_1, c_2 > 0$ are constants depending polynomially only on the subgaussian moment.

Next, we present a probabilistic bound on the deviation of the norm of weighted sum of squared Gaussian random variables from its mean. This is a direct extension of [81, Theorem 5.2.2].

Lemma 4. *Let $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ be a Gaussian random vector and $\lambda_1, \dots, \lambda_d > 0$ be constants. It holds for any $t > 0$ that*

$$\mathbb{P}\left(\left|\sqrt{\sum_{i=1}^d \lambda_i^2 x_i^2} - \sqrt{\sum_{i=1}^d \lambda_i^2}\right| \geq t + 2\lambda_{\max}\right) \leq 2 \exp\left(-\frac{t^2}{2\lambda_{\max}^2}\right), \quad (49)$$

where $\lambda_{\max} = \max\{\lambda_i : i \in [d]\}$.

Based on the above lemma, we can further show the following concentration inequalities to estimate the norm of the standard norm Gaussian random vector.

Lemma 5. *Suppose that $\mathbf{a}_i \stackrel{i.i.d.}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ is a Gaussian random vector for each $i \in [N]$. The following statements hold:*

(i) *It holds for all $i \in [N]$ with probability at least $1 - N^{-1}$ that*

$$\left|\|\mathbf{a}_i\| - \sqrt{d}\right| \leq 2\sqrt{\log N} + 2. \quad (50)$$

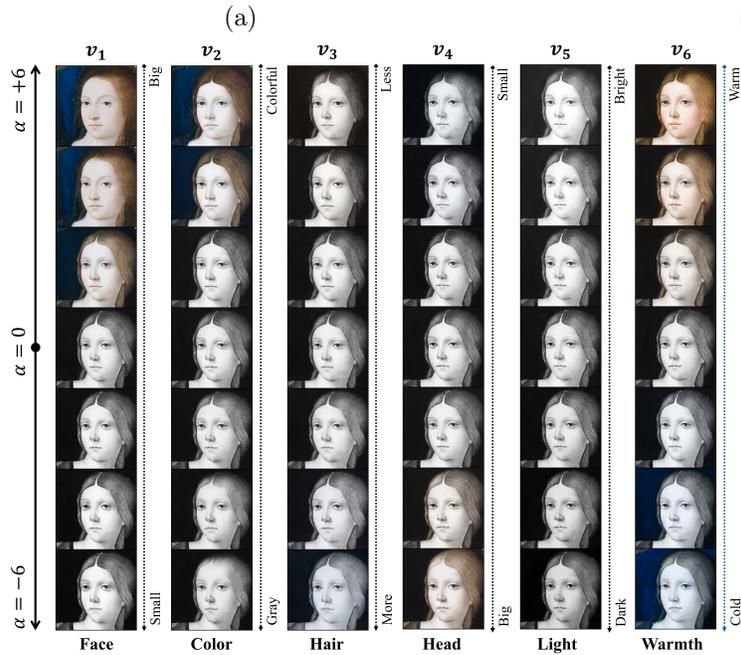
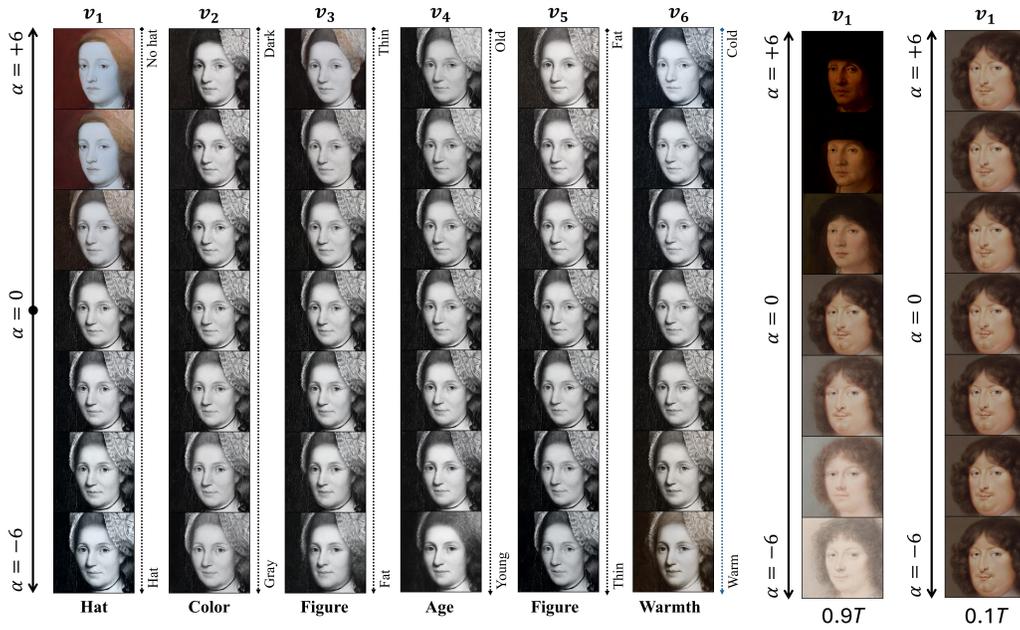
(ii) *Let $\mathbf{V} \in \mathcal{O}^{n \times d}$ be given. For all $i \in C_k^*$ and all $k \in [K]$, it holds with probability at least $1 - 2N^{-1}$ that*

$$\left|\|\mathbf{V}^T \mathbf{U}_k^* \mathbf{a}_i\| - \|\mathbf{V}^T \mathbf{U}_k^*\|_F\right| \leq 2\sqrt{\log N} + 2. \quad (51)$$

Proof. (i) Applying Lemma 4 to $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$, together with setting $t = 2\sqrt{\log N}$ and $\lambda_j = 1$ for all $j \in [d]$, yields

$$\mathbb{P}\left(\left|\|\mathbf{a}_i\| - \sqrt{d}\right| \geq 2\sqrt{\log N} + 2\right) \leq 2N^{-2}.$$

This, together with the union bound, yields that (50) holds with probability $1 - N^{-1}$.



(c)

Figure 9: Correspondence between the singular vectors of the Jacobian of the DAE and semantic image attributes. (a,c) Additional examples when $t = 0.7T$. (b) Ablation studies when $t = 0.1T$ and $0.9T$.

(ii) Let $\mathbf{V}^T \mathbf{U}_k^* = \mathbf{P} \mathbf{\Sigma} \mathbf{Q}^T$ be a singular value decomposition of $\mathbf{V}^T \mathbf{U}_k^*$, where $\mathbf{\Sigma} \in \mathbb{R}^{d \times d}$ with the diagonal elements $0 \leq \sigma_d \leq \dots \leq \sigma_1 \leq 1$ being the singular values of $\mathbf{V}^T \mathbf{U}_k^*$ and $\mathbf{P}, \mathbf{Q} \in \mathcal{O}^d$. This, together with the orthogonal invariance of the Gaussian distribution, yields

$$\|\mathbf{V}^T \mathbf{U}_k^* \mathbf{a}_i\| = \|\mathbf{\Sigma} \mathbf{Q}^T \mathbf{a}_i\| \stackrel{d}{=} \|\mathbf{\Sigma} \mathbf{a}_i\| = \sqrt{\sum_{j=1}^d \sigma_j^2 a_{ij}^2}. \quad (52)$$

Using Lemma 4 with setting $t = 2\sigma_1 \sqrt{\log N}$ and $\lambda_j = \sigma_j \leq 1$ for all j yields

$$\mathbb{P}\left(\left|\|\mathbf{V}^T \mathbf{U}_k^* \mathbf{a}_i\| - \|\mathbf{V}^T \mathbf{U}_k^*\|_F\right| \geq \sigma_1 \alpha\right) = \mathbb{P}\left(\left|\sqrt{\sum_{j=1}^d \sigma_j^2 a_{ij}^2} - \sqrt{\sum_{j=1}^d \sigma_j^2}\right| \geq \sigma_1 \alpha\right) \leq 2N^{-2}.$$

This, together with $\sigma_1 \leq 1$ and the union bound, yields (51). \square

Next, We present a spectral bound on the covariance estimation for the random vectors generated by the normal distribution.

Lemma 6. *Suppose that $\mathbf{a}_1, \dots, \mathbf{a}_N \in \mathbb{R}^d$ are i.i.d. standard normal random vectors, i.e., $\mathbf{a}_i \stackrel{i.i.d.}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ for all $i \in [N]$. Then, it holds with probability at least $1 - 2N^{-2}$ that*

$$\left\|\frac{1}{N} \sum_{i=1}^N \mathbf{a}_i \mathbf{a}_i^T - \mathbf{I}_d\right\| \leq \frac{9(\sqrt{d} + \sqrt{\log N})}{\sqrt{N}}, \quad (53)$$

Proof. According to [81, Theorem 4.7.1], it holds that

$$\mathbb{P}\left(\left\|\frac{1}{N} \sum_{i=1}^N \mathbf{a}_i \mathbf{a}_i^T - \mathbf{I}_d\right\| \geq \frac{9(\sqrt{d} + \eta)}{\sqrt{N}}\right) \leq 2 \exp(-2\eta^2),$$

where $\eta > 0$. Plugging $\eta = \sqrt{\log N}$ into the above inequality yields

$$\mathbb{P}\left(\left\|\frac{1}{N} \sum_{i=1}^N \mathbf{a}_i \mathbf{a}_i^T - \mathbf{I}_d\right\| \geq \frac{9(\sqrt{d} + \sqrt{\log N})}{\sqrt{N}}\right) \leq 2N^{-2}.$$

This directly implies (53). \square

Lemma 7. *Let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$ be positive semi-definite matrices. Then, it holds that*

$$\langle \mathbf{A}, \mathbf{B} \rangle \geq \lambda_{\min}(\mathbf{A}) \text{Tr}(\mathbf{B}). \quad (54)$$

Proof. Let $\mathbf{U} \mathbf{\Lambda} \mathbf{U}^T = \mathbf{A}$ be an eigenvalue decomposition of \mathbf{A} , where $\mathbf{U} \in \mathcal{O}^n$ and $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$ is a diagonal matrix with diagonal entries $\lambda_1 \geq \dots \geq \lambda_n \geq 0$ being the eigenvalues. Then, we compute

$$\langle \mathbf{A}, \mathbf{B} \rangle = \langle \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T, \mathbf{B} \rangle = \langle \mathbf{\Lambda}, \mathbf{U} \mathbf{B} \mathbf{U}^T \rangle \geq \lambda_{\min}(\mathbf{A}) \text{Tr}(\mathbf{U} \mathbf{B} \mathbf{U}^T) = \lambda_{\min}(\mathbf{A}) \text{Tr}(\mathbf{B}),$$

where the inequality follows from $\lambda_i \geq 0$ for all $i \in [N]$ and \mathbf{B} is a positive semidefinite matrix. \square