

---

# On the Generalization Properties of Diffusion Models

---

**Puheng Li\*\***

Department of Statistics  
Stanford University  
puhengli@stanford.edu

**Zhong Li\*†**

Machine Learning Group  
Microsoft Research Asia  
lzhong@microsoft.com

**Huishuai Zhang†**

Machine Learning Group  
Microsoft Research Asia  
huzhang@microsoft.com

**Jiang Bian**

Machine Learning Group  
Microsoft Research Asia  
jiabia@microsoft.com

## Abstract

Diffusion models are a class of generative models that serve to establish a stochastic transport map between an empirically observed, yet unknown, target distribution and a known prior. Despite their remarkable success in real-world applications, a theoretical understanding of their generalization capabilities remains underdeveloped. This work embarks on a comprehensive theoretical exploration of the generalization attributes of diffusion models. We establish theoretical estimates of the generalization gap that evolves in tandem with the training dynamics of score-based diffusion models, suggesting a polynomially small generalization error ( $O(n^{-2/5} + m^{-4/5})$ ) on both the sample size  $n$  and the model capacity  $m$ , evading the curse of dimensionality (i.e., not exponentially large in the data dimension) when *early-stopped*. Furthermore, we extend our quantitative analysis to a *data-dependent* scenario, wherein target distributions are portrayed as a succession of densities with progressively increasing distances between modes. This precisely elucidates the *adverse* effect of “*modes shift*” in ground truths on the model generalization. Moreover, these estimates are not solely theoretical constructs but have also been confirmed through numerical simulations. Our findings contribute to the rigorous understanding of diffusion models’ generalization properties and provide insights that may guide practical applications.

## 1 Introduction

As an emerging family of deep generative models, diffusion models (DMs; [16, 45]) have experienced a surge in popularity, owing to their unparalleled performance in a wide range of applications ([22, 39, 68, 15, 3, 35, 4, 60, 33, 69, 64, 63, 53]). This has led to notable commercial successes, such as DALL·E ([34]), Imagen ([38]), and Stable Diffusion ([36]). Mathematically, diffusion models learn an unknown underlying distribution through a two-stage process: (i) first, successively and gradually injecting random noises (forward process); (ii) then reversing the forward process through denoising for sampling purposes (reverse process). To achieve this, an equivalent formulation of diffusion models called score-based generative models (SGMs; [50, 52]) is employed. SGMs implement the aforementioned two-stage process via the continuous dynamics represented by a joint group of coupled stochastic differential equations (SDEs) ([54, 20]).

---

\*This work was done when Puheng Li was an undergraduate student at Peking University and a research intern at MSRA. The first two authors contributed equally to this work and are ordered alphabetically.

†Corresponding author.

Despite their impressive empirical performance, the theoretical foundation of DMs/SGMs remains underexplored. Generally, fundamental theoretical questions can be categorized into several aspects. By considering machine learning models as mathematical function classes from certain spaces, one can identify three central aspects: approximation, optimization and generalization. At the forefront lies the generalization problem, which aims to characterize the learning error between the learned and ground truth distributions.

The development of generalization theory for diffusion models is pressing due to both theoretical and practical concerns:

- In theory, the generalization issues of generative modeling (or learning for distributions) may exhibit as the memorization phenomenon, if the modeled distribution is eventually trained to converge to the empirical distribution only associated with training samples. Intuitively, memorization arises from two reasons: (i) it is useful for the hypothesis space to be large enough to approximate highly complex underlying target distributions (universal convergence; [65]); (ii) the underlying distribution is unknown in practice, and one can only use a dataset with finite samples drawn from the target distribution. Rigorous mathematical characterizations of memorization are developed for bias potential models and GANs in [66] and [67], respectively. A natural question is, does a similar phenomenon occur for diffusion models? To answer this, a thorough investigation of generalization properties for DMs/SGMs is required.
- In practice, the generalization capability of diffusion models is also an essential requirement, as the memorization can lead to potential privacy and copyright risks when models are deployed. Similar to other generative models and large language models (LLMs) [5, 70, 19, 6], diffusion models can also memorize and leak training samples [5, 46], hence can be subsequently attacked using specific procedures and algorithms [28, 18, 62]. Although there are defense methods developed to meet privacy and copyright standards ([11, 14, 58]), these approaches are often heuristic, without providing sufficient quantitative understandings particularly on diffusion models. Therefore, a comprehensive investigation of the generalization foundation of diffusion models, including both theoretical and empirical aspects, is of utmost importance in improving principled tutorial guidance in practice.

The current work develops the generalization theory of diffusion models in a mathematically rigorous manner. Our main results include the following:

- We derive an upper bound of the generalization gap for diffusion models along the training dynamics. This result suggests, with *early-stopping*, the generalization error of diffusion models scales polynomially small on the sample size ( $O(n^{-2/5})$ ) and the model capacity ( $O(m^{-4/5})$ ). Notably, the generalization error also escapes from the curse of dimensionality.
- This “uniform” bound is further extended to a *data-dependent* setting, where a sequence of unidimensional Gaussian mixtures distributions with an increasing modes’ distance is considered as the ground truth. This result characterizes the effect of “*modes shift*” quantitatively, which implies that the generalization capability of diffusion models is *adversely* affected by the distance between high-density regions of target distributions.
- The theoretical findings are numerically verified in simulations.

The rest of this paper is organized as follows. In Section 2, we discuss the related work on the convergence and training fronts of diffusion models and also the generalization aspects of other generative modeling methods. Section 3 is the central part, which includes the problem formulation, main results, and consequences. Section 4 includes numerical verifications on synthetic and real-world datasets<sup>3</sup>. All the details of proofs and experiments are found in the appendices.

## 2 Related Work

We review the related work on diffusion models concerning the central results in this paper.

- First, on the convergence theory, [7, 9] established elaborate error estimates between the modeled and target distribution given the discretization and time-dependent score matching

---

<sup>3</sup>Code is available at [https://github.com/lphLeo/Diffusion\\_Generalization](https://github.com/lphLeo/Diffusion_Generalization)



**Forward perturbation.** We start with the setting of unsupervised learning. Given an unlabeled dataset  $\mathcal{D}_x = \{\mathbf{x}_i\}_{i=1}^n \subset \mathbb{R}^d$  with the sample  $\mathbf{x}_i \stackrel{\text{i.i.d.}}{\sim} p_0(\mathbf{x})$ , where  $p_0$  denotes the underlying (ground truth or target) distribution, the forward diffusion process is defined as

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, t)dt + g(t)d\mathbf{W}_t, \quad \mathbf{x}(0) \sim p_0. \quad (1)$$

Here, the drift coefficient  $\mathbf{f}(\cdot, t) : \mathbb{R}^d \mapsto \mathbb{R}^d$  is a time-dependent vector-valued function, and the diffusion coefficient  $g(\cdot) : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}$  is a scalar function, and  $\mathbf{W}_t$  denotes the standard Wiener process (a.k.a., Brownian motion). The SDE (1) has a unique, strong solution under certain regularity conditions (i.e., globally Lipschitz coefficients in both state and time; see [31]). From now on, we denote by  $p_t(\mathbf{x}(t))$  the marginal distribution of  $\mathbf{x}(t)$ , and let  $p_{t|s}(\mathbf{x}(t)|\mathbf{x}(s))$  be the (perturbation) transition kernel from  $\mathbf{x}(s)$  to  $\mathbf{x}(t)$ ,  $0 \leq s < t \leq T < \infty$  with  $T$  as the time horizon. By appropriately selecting  $\mathbf{f}$  and  $g$ , one can force the SDE (1) to converge to a prior distribution (typically a Gaussian). Common examples include the (time-rescaled) Ornstein–Uhlenbeck (OU) process,<sup>4</sup> which is a special case of the linear version of (1):

$$d\mathbf{x} = f(t)\mathbf{x}dt + g(t)d\mathbf{W}_t, \quad \mathbf{x}(0) \sim p_0. \quad (2)$$

**Reverse sampling.** According to [1] and [54], both the following reverse-time SDE and probability flow ODE share the same marginal distribution as the forward-time SDE (1):

$$d\mathbf{x} = \left[ \mathbf{f}(\mathbf{x}, t) - g^2(t)\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) \right] dt + g(t)d\bar{\mathbf{W}}_t, \quad (3)$$

$$d\mathbf{x} = \left[ \mathbf{f}(\mathbf{x}, t) - \frac{1}{2}g^2(t)\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) \right] dt, \quad (4)$$

where  $\bar{\mathbf{W}}_t$  is a standard Wiener process when time flows backwards from  $T$  to 0, and  $dt$  is an infinitesimal negative time step. With the initial condition  $\mathbf{x}(T) \sim p_T \approx \pi$ , where  $\pi$  is a known prior distribution such as the Gaussian noise, one can (numerically) solve (3) or (4) to transform noises into samples from  $p_0$ , which is exactly the goal of generative modeling.

**Loss objectives.** The only remaining task is to estimate the unknown (Stein) score function  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ . This is achieved by minimizing the following weighted sum of denoising score matching ([57]) objectives:

$$\mathcal{L}(\boldsymbol{\theta}; \lambda(\cdot)) := \mathbb{E}_{t \sim \mathcal{U}(0, T)} \left[ \lambda(t) \cdot \mathbb{E}_{\mathbf{x}(0) \sim p_0} \left[ \mathbb{E}_{\mathbf{x}(t) \sim p_{t|0}} \left[ \|\mathbf{s}_{t, \boldsymbol{\theta}}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0))\|_2^2 \right] \right] \right], \quad (5)$$

with  $\boldsymbol{\theta}^* := \arg \min_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}; \lambda(\cdot))$ , where  $\mathcal{U}(0, T)$  denotes the uniform distribution over  $[0, T]$ , and  $\lambda(t) : [0, T] \mapsto \mathbb{R}_+$  is a weighting function, which is typically selected as

$$\lambda(t) \propto 1/\sqrt{\mathbb{E}_{\mathbf{x}(t) \sim p_{t|0}} \left[ \|\nabla_{\mathbf{x}(t)} \log p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0))\|_2^2 \right]} \quad (6)$$

according to ([54]). The score function  $\mathbf{s}_{t, \boldsymbol{\theta}} : \mathbb{R}^d \mapsto \mathbb{R}^d$  is time-dependent and can be parameterized as a neural network (encoded with the time information) such as the U-net([37]) architecture commonly applied in the field of image segmentation. Alternatively, one can also define the time-dependent score matching loss

$$\tilde{\mathcal{L}}(\boldsymbol{\theta}; \lambda(\cdot)) := \mathbb{E}_{t \sim \mathcal{U}(0, T)} \left[ \lambda(t) \cdot \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \|\mathbf{s}_{t, \boldsymbol{\theta}}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t))\|_2^2 \right] \right], \quad (7)$$

which is equivalent to (5) up to a constant independent of  $\boldsymbol{\theta}$  by [57, 49].

In practice, expectations in the objective (5) can be respectively estimated with empirical means over time steps in  $[0, T]$ , data samples from  $p_0$  and  $p_{t|0}$ , which is efficient when the drift coefficient  $\mathbf{f}(\cdot, t)$  is linear. Specifically, if the forward-time SDE takes the form of (2), the transition kernel  $p_{t|0}$  has a closed form ([40])

$$p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0)) = \mathcal{N}(\mathbf{x}(t); r(t)\mathbf{x}(0), r^2(t)v^2(t)\mathbf{I}_d), \quad (8)$$

where  $\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$  denotes the multivariate Gaussian distribution evaluated at  $\mathbf{x}$  with the expectation  $\boldsymbol{\mu}$  and covariance  $\boldsymbol{\Sigma}$ , and  $r(t) := e^{\int_0^t f(\zeta)d\zeta}$ ,  $v(t) := \sqrt{\int_0^t \frac{g^2(\zeta)}{r^2(\zeta)} d\zeta}$ .

<sup>4</sup>The OU process is the unique time-homogeneous Markov process which is also a Gaussian process, with the stationary distribution as the standard Gaussian distribution.

**Training.** We aim to investigate the gradient flow training dynamics over the empirical loss

$$\frac{d}{d\tau} \hat{\boldsymbol{\theta}}_n(\tau) = -\nabla_{\hat{\boldsymbol{\theta}}_n(\tau)} \hat{\mathcal{L}}_n(\hat{\boldsymbol{\theta}}_n(\tau); \lambda(\cdot)), \quad \hat{\boldsymbol{\theta}}_n(0) := \hat{\boldsymbol{\theta}}_n^0, \quad (9)$$

where  $\hat{\mathcal{L}}_n$  is the Monte-Carlo estimation of  $\mathcal{L}$  defined in (5) on the training dataset,<sup>5</sup> with an auxiliary gradient flow over the population loss

$$\frac{d}{d\tau} \boldsymbol{\theta}(\tau) = -\nabla_{\boldsymbol{\theta}(\tau)} \mathcal{L}(\boldsymbol{\theta}(\tau); \lambda(\cdot)), \quad \boldsymbol{\theta}(0) := \boldsymbol{\theta}^0 = \hat{\boldsymbol{\theta}}_n^0. \quad (10)$$

In both cases, the weighting function  $\lambda(\cdot)$  is selected as in (6). Denote the score function learned at the training time  $\tau$  evaluated at the SDE time  $t$  with respect to the empirical loss and population loss as  $\mathbf{s}_{t, \hat{\boldsymbol{\theta}}_n(\tau)}(\mathbf{x}(t))$  and  $\mathbf{s}_{t, \boldsymbol{\theta}(\tau)}(\mathbf{x}(t))$ , respectively. The corresponding density functions, denoted by  $p_{t, \hat{\boldsymbol{\theta}}_n(\tau)}(\mathbf{x}(t))$  and  $p_{t, \boldsymbol{\theta}(\tau)}(\mathbf{x}(t))$ , are obtained by solving

$$\nabla_{\mathbf{x}(t)} \log p_{t, \hat{\boldsymbol{\theta}}_n(\tau)}(\mathbf{x}(t)) = \mathbf{s}_{t, \hat{\boldsymbol{\theta}}_n(\tau)}(\mathbf{x}(t)), \quad \nabla_{\mathbf{x}(t)} \log p_{t, \boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) = \mathbf{s}_{t, \boldsymbol{\theta}(\tau)}(\mathbf{x}(t)), \quad (11)$$

and then normalizing, respectively.

**Score networks.** We parameterize the score function  $\mathbf{s}_{t, \boldsymbol{\theta}}$  as the following random feature model

$$\mathbf{s}_{t, \boldsymbol{\theta}}(\mathbf{x}) := \frac{1}{m} \mathbf{A} \sigma(\mathbf{W} \mathbf{x} + \mathbf{U} e(t)) = \frac{1}{m} \sum_{i=1}^m \mathbf{a}_i \sigma(\mathbf{w}_i^\top \mathbf{x} + \mathbf{u}_i^\top e(t)), \quad (12)$$

where  $\sigma$  is the ReLU activation function,  $\mathbf{A} = (\mathbf{a}_1, \dots, \mathbf{a}_m) \in \mathbb{R}^{d \times m}$  is the *trainable* parameter, while  $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_m)^\top \in \mathbb{R}^{m \times d}$  and  $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_m)^\top \in \mathbb{R}^{m \times d_e}$  are randomly initialized and *frozen* during training, and  $e : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}^{d_e}$  is the embedding function concerning the time information. Assume that  $\mathbf{a}_i$ ,  $\mathbf{w}_i$  and  $\mathbf{u}_i$  are i.i.d. sampled from an underlying distribution  $\rho$ . Then, as  $m \rightarrow \infty$ , we get

$$\begin{aligned} \mathbf{s}_{t, \boldsymbol{\theta}}(\mathbf{x}) &\rightarrow \bar{\mathbf{s}}_{t, \bar{\boldsymbol{\theta}}}(\mathbf{x}) := \mathbb{E}_{(\mathbf{a}, \mathbf{w}, \mathbf{u}) \sim \rho} [\mathbf{a} \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top e(t))] \\ &= \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [\mathbf{a}(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top e(t))], \end{aligned} \quad (13)$$

with  $\mathbf{a}(\mathbf{w}, \mathbf{u}) := \frac{1}{\rho_0(\mathbf{w}, \mathbf{u})} \int_{\mathbb{R}^d} \mathbf{a} \rho(\mathbf{a}, \mathbf{w}, \mathbf{u}) d\mathbf{a}$  and  $\rho_0(\mathbf{w}, \mathbf{u}) := \int_{\mathbb{R}^d} \rho(\mathbf{a}, \mathbf{w}, \mathbf{u}) d\mathbf{a}$ . By the positive homogeneity property of the ReLU activation, we can assume that  $\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1$  w.l.o.g.

One can view  $\bar{\mathbf{s}}_{t, \bar{\boldsymbol{\theta}}}(\mathbf{x})$  as a continuous version of the random feature model. Correspondingly, the optimal solution is denoted as  $\bar{\boldsymbol{\theta}}^*$  when replacing the parameterized score function  $\mathbf{s}_{t, \boldsymbol{\theta}}(\mathbf{x})$  in the loss objective (5) or (7) by  $\bar{\mathbf{s}}_{t, \bar{\boldsymbol{\theta}}}(\mathbf{x})$ . Define the kernel  $k_{\rho_0}(\mathbf{x}, \mathbf{x}') := \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [\sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top e(t)) \sigma(\mathbf{w}^\top \mathbf{x}' + \mathbf{u}^\top e(t))]$ , and let  $\mathcal{H}_{k_{\rho_0}}$  be the induced reproducing kernel Hilbert space (RKHS; [2]), we have  $\bar{\mathbf{s}}_{t, \bar{\boldsymbol{\theta}}} \in \mathcal{H}_{k_{\rho_0}}$  if the RKHS norm  $\|\bar{\mathbf{s}}_{t, \bar{\boldsymbol{\theta}}}\|_{\mathcal{H}_{k_{\rho_0}}}^2 := \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [\|\mathbf{a}(\mathbf{w}, \mathbf{u})\|_2^2] = \|\mathbf{a}\|_2^2_{L^2(\rho_0)} < \infty$ , and the corresponding discrete version can be defined by the empirical average, i.e.,  $\|\mathbf{s}_{t, \boldsymbol{\theta}}\|_{\mathcal{H}_{k_{\rho_0}}}^2 := \frac{1}{m} \|\mathbf{A}\|_F^2 = \frac{1}{m} \sum_{i=1}^m \|\mathbf{a}(\mathbf{w}_i, \mathbf{u}_i)\|_2^2$ .

**Remark 1.** *There are more modern and complex mathematical tools such as neural tangent kernels (NTKs) and mean fields that can be selected as the score networks. Employing these modern tools is valuable at least for theoretical completeness and we leave these as the future work.*<sup>6</sup>

The goal is to measure and bound the generalization error evolving with the gradient flow training dynamics (9) between the learned distribution and target distribution, using the common Kullback–Leibler (KL) divergence.

**Definition 1** (KL divergence). *Given two distributions  $p$  and  $q$ , the KL divergence from  $q$  to  $p$  is defined as  $D_{\text{KL}}(p||q) = \int_{\mathbb{R}^d} p(\mathbf{x}) \log \left( \frac{p(\mathbf{x})}{q(\mathbf{x})} \right) d\mathbf{x}$ .*

Based on the above definitions, the generalization gap along the gradient flow training dynamics (9) is formulated as  $D_{\text{KL}} \left( p_0 || p_{0, \hat{\boldsymbol{\theta}}_n(\tau)} \right)$ , which is only a function of the training time  $\tau$ . The goal is to estimate  $D_{\text{KL}} \left( p_0 || p_{0, \hat{\boldsymbol{\theta}}_n(\tau)} \right)$ .

<sup>5</sup>Recall the dataset  $\mathcal{D}_{\mathbf{x}} = \{\mathbf{x}_i\}_{i=1}^n \subset \mathbb{R}^d$ , and we take  $\mathcal{D}_t = \{t_i\}_{i=1}^n \subset [0, T]$  for convenience.

<sup>6</sup>For example, the extension to NTKs is possible, since the training of random feature models follows a specific NTK regime (only the last layer is updated) in the output space (instead of the parameter space), which can be properly analyzed in the infinite-width regime.

### 3.2 Main Results

In this section, we state the main results of the generalization capability of the (score-based) diffusion models and how it evolves as the training proceeds. Based on the formulation in Section 3.1, we theoretically derive several upper bounds to estimate  $D_{\text{KL}}(p_0 \| p_{0, \hat{\theta}_n(\tau)})$  under different settings. The results cover both the positive and negative aspects: in the data-independent setting where the target distribution has finite support, the generalization gap is proved to be small; while in the data-dependent setting where the target distribution possesses shift modes, the generalization is adversely affected by the modes' distance.

#### 3.2.1 Data-Independent Generalization Gap

In this section, we provide the characterization of the generalization capability for diffusion models given a target distribution defined on a finite domain. Generally, the KL divergence from the learned distribution  $p_{0, \hat{\theta}_n(\tau)}$  at the training time  $\tau$  to the target distribution  $p_0$  can be estimated as follows.

**Theorem 1.** *Suppose that the target distribution  $p_0$  is continuously differentiable and has a compact support set, i.e.,  $\|\mathbf{x}\|_\infty$  is uniformly bounded, and there exists a reproducing kernel Hilbert space (RKHS)  $\mathcal{H} (:= \mathcal{H}_{k_{p_0}})$  such that  $\bar{\mathbf{s}}_{0, \bar{\theta}^*} \in \mathcal{H}$ . Assume that the initial loss, trainable parameters, the embedding function  $e(t)$  and weighting function  $\lambda(t)$  are all bounded. Then for any  $\delta > 0$ ,  $\delta \ll 1$ , with the probability of at least  $1 - \delta$ , we have*

$$D_{\text{KL}}(p_0 \| p_{0, \hat{\theta}_n(\tau)}) \lesssim \left[ \frac{\tau^4}{mn} + \frac{\tau^3}{m^2} + \frac{1}{\tau} \right] + \left[ \frac{1}{m} + \bar{\mathcal{L}}(\bar{\theta}^*) + \tilde{\mathcal{L}}(\theta^*) \right] + D_{\text{KL}}(p_T \| \pi), \quad \tau \geq 1,$$

where  $\lesssim$  hides the term  $d \log(d + 1)$ , the polynomials of  $\log(1/\delta^2)$ , finite RKHS norms and universal positive constants only depending on  $T$ .

**Remark 2.** *Since  $p_0$  is compactly supported, the target score function  $\mathbf{s}_0(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_0(\mathbf{x})$  is also defined on a compact domain. According to [12, 13],  $\mathbf{s}_0$  is contained in the Barron function space with a finite Barron norm, and hence in a certain RKHS with a finite RKHS norm. Therefore, it is reasonable to require that the global minimizer  $\mathbf{s}_{0, \theta^*}$  or  $\bar{\mathbf{s}}_{0, \bar{\theta}^*}$  is also contained in some RKHS.*

**Proof sketch.** Theorem 1 is proved via the following procedure.

1. According to Theorem 1 in [49], the KL divergence on the left-hand side can be upper bounded by the population loss of the trained model up to a small error. That is,

$$D_{\text{KL}}(p_0 \| p_{0, \hat{\theta}_n(\tau)}) \leq \tilde{\mathcal{L}}(\hat{\theta}_n(\tau); g^2(\cdot)) + D_{\text{KL}}(p_T \| \pi). \quad (14)$$

2. We use the model trained with respect to the population loss (10) to perform the decomposition:

$$\begin{aligned} \tilde{\mathcal{L}}(\hat{\theta}_n(\tau)) &= \left[ \tilde{\mathcal{L}}(\hat{\theta}_n(\tau)) - \tilde{\mathcal{L}}(\theta(\tau)) \right] + \tilde{\mathcal{L}}(\theta(\tau)) \\ &\lesssim \left[ \tilde{\mathcal{L}}(\hat{\theta}_n(\tau)) - \tilde{\mathcal{L}}(\theta(\tau)) \right] + \bar{\mathcal{L}}(\bar{\theta}(\tau)) + \text{Monte Carlo} \triangleq I_1 + I_2 + I_3, \end{aligned} \quad (15)$$

where we omit the weighting function  $g^2(\cdot)$  for simplicity. Here,  $\bar{\mathcal{L}}$  is the loss objective obtained by replacing  $\mathbf{s}_{t, \theta}$  in  $\tilde{\mathcal{L}}$  (defined in (7)) by  $\bar{\mathbf{s}}_{t, \bar{\theta}}$ , and  $I_3$  summarizes the resulting Monte Carlo error.

3.  $I_3$  can be estimated via a similar argument as in [24] (Lemma 48).
4.  $I_2$  can be upper bounded via a standard analysis on the gradient flow dynamics over convex objectives.
5.  $I_1$  can be reduced as the norm product of  $\mathbf{s}_{0, \hat{\theta}_n(\tau)}$ ,  $\mathbf{s}_{0, \theta(\tau)}$  and their gap, then
  - (a) the former can be bounded with a square-root rate growth via a general norm estimate of parameters trained under the gradient flow dynamics;
  - (b) the latter can be estimated by the Rademacher complexity (see e.g., Chapter 26 in [43]).

Combining all above gives the desired result. The detailed proof is found in Appendix A.1.

**Discussion on error bounds.** The three error terms are further analyzed as follows.

- The first term is the main error, which implies an *early-stopping* generalization gap. In fact, if one selects an early-stopping time  $\tau_{\text{es}}$  as  $\tau_{\text{es}} = \Theta\left(n^{\frac{2}{5}}\right)$ , and let  $m \sim n$ , we have

$$D_{\text{KL}}\left(p_0 \| p_{0, \hat{\theta}_n(\tau_{\text{es}})}\right) \lesssim (1/n)^{\frac{2}{5}} + (1/m)^{\frac{4}{5}}. \quad (16)$$

- The second term is  $m$ -dependent and corresponds to the approximation error, which is  $o(1)$  when  $m \gg 1$ . In fact, the random feature model is a universal approximator to Lipschitz continuous functions on a compact domain (Theorem 6 in [17]). See more details in Appendix A.1 (the last paragraph).
- The third term is exponentially small in  $T$  since  $\pi$  (e.g. the Gaussian density) is log-Sobolev, according to a classical result in e.g. [55] (Theorem 3.20, Theorem 3.24 and Remark 3.26).

**Remark 3.** In practice, it is common to use the test error to evaluate the generalization performance. For diffusion models, a straightforward approach is to compute the negative log-likelihood (averaged in bits/dim; equivalent to the KL divergence) on the test dataset during training with the instantaneous change-of-variable formula ([8]) and probability flow ODE (defined in (4)), where the true score function  $\nabla_x \log p_t(x)$  is replaced by  $s_{t, \theta(\tau)}(x)$ .

**Remark 4.** Previous literature has established similar bounds for bias potential models ([66]) and GANs ([67]). Theorem 1 extends the corresponding results to the setting of diffusion models. Furthermore, this upper bound is finer in the sense that it incorporates the information regarding the model capacity (the hidden dimension  $m$  in this case), which shows that more parameters benefit the learning and generalization, as expected.

### 3.2.2 Data-Dependent Generalization Gap

In Section 3.2.1, we derive estimates on the generalization error for diffusion models along the training dynamics, where the target distribution is assumed to be finitely supported. In reality, this is often not the case, where target distributions usually possess distant multi-modes, from simple Gaussian mixtures to complicated Boltzmann distributions of physical systems ([29, 30]). Under these settings, the above analysis in Section 3.2.1 can not directly apply since the data domain is unbounded. It remains a problem to quantitatively characterize the generalization behavior of diffusion models given these target distributions with distant multi-modes or modes shift.

To provide a fine-grained demonstration of the generalization capability of diffusion models when applied to learn distributions with distant multi-modes, as an illustrating example, the Gaussian mixture with two modes is selected as the target distribution.

**Theorem 2.** Suppose the target distribution  $p_0$  is a one-dimensional 2-mode Gaussian mixture:  $p_0(x) = q_1 \mathcal{N}(x; -\mu, 1) + q_2 \mathcal{N}(x; \mu, 1)$ , where  $\mu > \sqrt{\log(1/\delta^2)}$ ,  $q_1, q_2 > 0$  with  $q_1 + q_2 = 1$  are all constants. Under the conditions of Theorem 1 (except the uniform boundness of inputs), we have

$$D_{\text{KL}}\left(p_0 \| p_{0, \hat{\theta}_n(\tau)}\right) \lesssim \text{Poly}(\mu) \left[ \frac{\tau^4}{mn} + \frac{\tau^3}{m^2} \right] + \frac{1}{\tau} + \left[ \frac{\mu^2}{m} + \tilde{\mathcal{L}}(\bar{\theta}^*) + \tilde{\mathcal{L}}(\theta^*) \right] + D_{\text{KL}}(p_T \| \pi),$$

where  $\tau \geq 1$ ,  $\lesssim$  hides the polynomials of  $\log(1/\delta^2)$ , finite RKHS norms and universal positive constants only depending on  $T$ .

**Proof sketch.** Theorem 2 is proved following a similar procedure with Theorem 1, except the input data  $x$  does not have a uniform bound here. This problem mainly affects the last step (5 (b)) in the proof sketch of Theorem 1, and can be handled by using the fact

$$|x| \in [\mu - \sqrt{\log(1/\delta^2)}, \mu + \sqrt{\log(1/\delta^2)}] = \Theta(\mu) \quad (17)$$

given the target Gaussian mixture distribution. The detailed proof is found in Appendix A.2.

**Remark 5.** Theorem 2 indicates that, even for a simple target distribution (e.g. a one-dimensional 2-mode Gaussian mixture), the generalization error of diffusion models can be polynomially large regarding the modes' distance. Although Theorem 2 provides only an upper bound, the modes shift effect holds due to the model-target inconsistency (the last paragraph in Appendix A.2) and the following consistent experiments (Section 4.1.2).

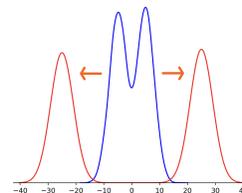


Figure 2: An illustration of modes shift.

## 4 Numerical Verifications

In this section, we numerically verify the previous theoretical results and insights (early-stopping generalization and modes shift effect) on both synthetic datasets and real-world datasets.

### 4.1 Simulations on Synthetic Datasets

#### 4.1.1 Early-Stopping Generalization

First, we illustrate the early-stopping generalization gap established in Theorem 1. We select the one-hidden-layer neural network with Swish activations as the score network, which is trained using the SGD optimizer with a fixed learning rate 0.5. The target distribution is set to be a one-dimensional 2-mode Gaussian mixture with the modes' distance equalling 6, and the number of data samples is 1000. We measure the KL divergence from the trained model to the target distribution  $D_{\text{KL}}(p_0 \| p_{0, \hat{\theta}_n(\tau)})$  along with the training epochs.

From Figure 3, one can observe that the KL divergence achieves its minimum at approximately the 800-th training epoch, and it starts to increase after this turning point. The experimental results are consistent with Theorem 1 (over multiple runs), which states that there exist early-stopping times when diffusion models can generalize well, indicating the effectiveness of our upper bound. Further, the KL divergence begins to oscillate after the minimum point, which may suggest a phase transition in the training dynamics, and the transition point is around the (optimal) early-stopping time.

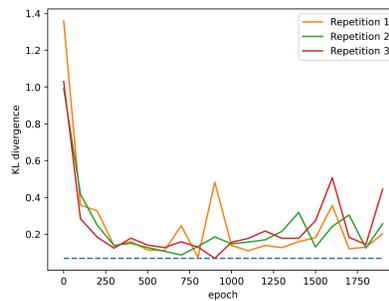


Figure 3: The KL divergence dynamics.

#### 4.1.2 Modes Shift Effect

Next, we numerically test the relationship between the modes' distance and generalization (density estimation) performance. All the configurations remain the same as Section 4.1.1, except that the target Gaussian mixtures have different modes' distances.

In Figure 4, the modeled distributions exhibit the following two-stage dynamics: (i) first gradually fitting the two modes (epoch = 100  $\rightarrow$  1000); (ii) then diverging (epoch = 1000  $\rightarrow$  1900). This aligns with the KL divergence dynamics (Figure 3) and again verifies the corresponding theoretical results (Theorem 1). However, Figure 5 shows that when the modes are distant from each other, there is *difficulty* in the learning process. In Figure 5, the optimal generalization is achieved at epoch = 100, but is still far from well generalizing. As the training proceeds, the learned model is almost always a single-mode distribution (epoch = 1000, 1900). This phenomenon is aligned with the results established in Theorem 2, which states that when there are distant modes in the target distribution, the generalization performance is relatively poor.

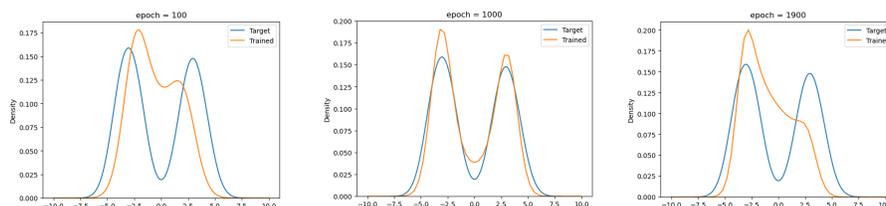


Figure 4: The training dynamics when the distance between two modes is 6 ( $\mu = 3$ ).

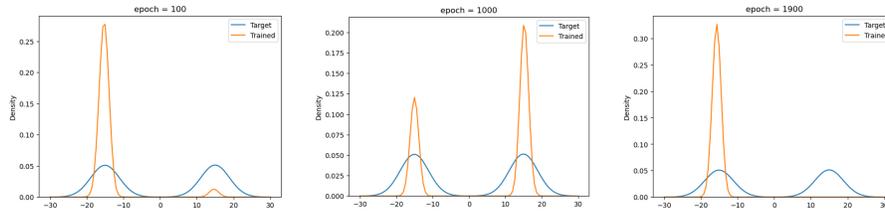


Figure 5: The training dynamics when the distance between two modes is 30 ( $\mu = 15$ ).

## 4.2 Simulations on Real-World Datasets

In this subsection, we verify our results on the MNIST dataset using the standard U-net architecture as the score network, which suggests that the adverse effect of modes shift on the generalization performance of diffusion models also appears *in general*.

The setup is as follows. First, we perform a  $K$ -means clustering on  $\mathcal{D}$  ( $\mathcal{D}$  denote the MNIST dataset) to get  $\mathcal{D} = \bigcup_{k=1}^K \mathcal{D}_k$ , and  $\bar{x}_k$  as the center of  $\mathcal{D}_k$ ,  $k = 1, 2, \dots, K$ . Let  $(i^*, j^*) := \arg \max_{i \neq j} \|\bar{x}_i - \bar{x}_j\|_2$ , and  $\mathcal{D}_{\text{farthest}} := \mathcal{D}_{i^*} \cup \mathcal{D}_{j^*}$ .  $\mathcal{D}_{\text{nearest}}$  is similarly constructed by  $\arg \min$  indices. Then, by randomly selecting the same number of data samples and using the same configuration, we train two separate diffusion models on  $\mathcal{D}_{\text{farthest}}$  and  $\mathcal{D}_{\text{nearest}}$ , respectively, and then perform inference (sampling). The training loss curves and sampling results are shown in Figure 6 and Figure 7, respectively. One can observe a significant performance gap: the diffusion model trained on  $\mathcal{D}_{\text{farthest}}$  appears a higher learning loss and worse sampling quality compared to those of  $\mathcal{D}_{\text{nearest}}$ .

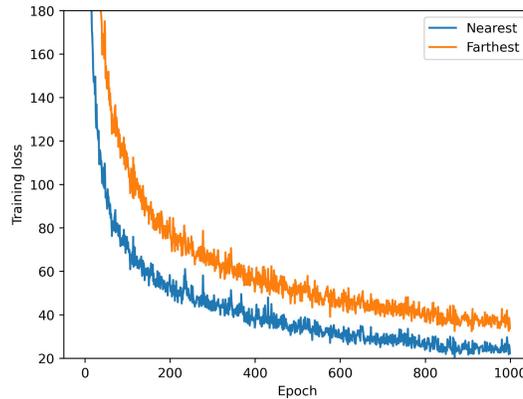


Figure 6: The training loss dynamics.



Figure 7: Sampling of the farthest (left) and nearest (right) clusters.

### 4.3 Discussion

We compare the results developed in this work with former corresponding literature as follows:

- The previous work [21] also studied the adverse effect of modes shift, which particularly reported a contrastive simulation indicating the degraded performance when modeling Gaussian mixtures with the increasing distance between modes (see Figure 2 in [21] and compare with Figure 4 and Figure 5). However, the results therein are established and tested under the “pure” score matching setting, without the denoising or time-dependent dynamics. As a comparison, Theorem 2 establishes a theoretical estimate on the generalization gap for diffusion models that requires *denoising*, and this upper bound *directly* depends on the distance between modes of target distributions, instead of a circuitous characterization in [21] to attribute the difficulty of learning modes shift to increased isoperimetry of corresponding target distributions.
- Similar adverse effect of modes shift has also been theoretically analyzed and numerically verified on recurrent neural networks (RNNs), see e.g., [23, 24]. There, the modes shift is understood as a type of long-term memory. This is the phenomenon of the “curse of memory”: When there is long-term memory in the target, it requires a large number of parameters for the approximation. Meanwhile, the training process will suffer from severe slowdowns. Both of these effects can be exponentially more pronounced with increasing memory or modes shift.
- The very recent work [42] also considered the problem of learning Gaussian mixtures using the denoising diffusion probabilistic model (DDPM) objective, but under a *teacher-student* setting. That is, given the Gaussian mixtures target, [42] parametrized the score network model in the same form of the score function target, with the goal to identify true parameters. Consequently, there are all positive convergence results developed in [42], despite that the time and sample complexity increases with the distance between modes. As a comparison, Theorem 2 adopts a pre-selected score network model without incorporating any information from the ground truth, and hence establish negative results concerning the modes shift. In fact, if the goal is to identify only true positions of the target Gaussian mixture (this is exactly the setting of [42]), the teacher-student setup seems not necessary (see Figure 5, where the true position is also learned efficiently using the one-hidden-layer Swish neural network as the score network, but the modes are always weighed incorrectly). In addition, [42] did not take the whole denoising dynamics into account. That is, the gradient descent (GD) analysis therein was performed on the denoising score matching objective successively at only two time stages: a larger  $t_1$  (“high noise”) and a smaller  $t_2$  (“low noise”), which is often not the case in practice.

## 5 Conclusion

In this paper, we provide a theoretical analysis of the fundamental generalization aspect of training diffusion models under both the data-independent and data-dependent settings, and early-stopping estimates of the generalization gap along the training dynamics are derived. Quantitatively, the data-independent results indicate a polynomially small generalization error that escapes from the curse of dimensionality, while the data-dependent results suggest the adverse effect of modes shift in target distributions. Numerical simulations have illustrated and verified these theoretical analyses. This work forms a basic starting point for understanding the intricacies of modern deep generative modeling and corresponding central concerns such as memorization, privacy, and copyright arising from practical applications and business products. More broadly, the approach here may have the potential to be extended to other variants in the diffusion models family, including a general SDE-based design space ([20]), consistency models ([48]), rectified flows ([27, 26]), Schrödinger bridges ([59, 56, 10, 44, 47, 25]), etc. These are certainly worthy of future exploration.

## Acknowledgements

We would like to thank Dr. Hongkang Yang for helpful discussions, and all the reviewers’ valuable feedback and insightful suggestions to improve this work.

## References

- [1] Brian D. O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- [2] Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, 1950.
- [3] Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021.
- [4] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18208–18218, 2022.
- [5] Nicholas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwal, Florian Tramer, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion models. In *USENIX Security Symposium*, pages 5253–5270, 2023.
- [6] Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramer, and Chiyuan Zhang. Quantifying memorization across neural language models. In *International Conference on Learning Representations*, 2023.
- [7] Hongrui Chen, Holden Lee, and Jianfeng Lu. Improved analysis of score-based generative modeling: User-friendly bounds under minimal smoothness assumptions. *International Conference on Machine Learning*, 202:4735–4763, 2023.
- [8] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K. Duvenaud. Neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 31:6571–6583, 2018.
- [9] Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru Zhang. Sampling is as easy as learning the score: Theory for diffusion models with minimal data assumptions. In *International Conference on Learning Representations*, 2023.
- [10] Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.
- [11] Tim Dockhorn, Tianshi Cao, Arash Vahdat, and Karsten Kreis. Differentially private diffusion models. *Transactions on Machine Learning Research*, 2023.
- [12] Weinan E, Chao Ma, and Lei Wu. A priori estimates of the population risk for two-layer neural networks. *Communications in Mathematical Sciences*, 17(5):1407–1425, 2019.
- [13] Weinan E, Chao Ma, and Lei Wu. The Barron space and the flow-induced function spaces for neural network models. *Constructive Approximation*, 55(1):369–406, 2022.
- [14] Sahra Ghalebikesabi, Leonard Berrada, Sven Gowal, Ira Ktena, Robert Stanforth, Jamie Hayes, Soham De, Samuel L Smith, Olivia Wiles, and Borja Balle. Differentially private diffusion models generate useful synthetic images. In *International Workshop on Trustworthy Federated Learning in Conjunction with IJCAI*, 2023.
- [15] Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and Lingpeng Kong. Diffuseq: Sequence to sequence text generation with diffusion models. In *International Conference on Learning Representations*, 2023.
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [17] Daniel Hsu, Clayton H. Sanford, Rocco Servedio, and Emmanouil Vasileios Vlatakis-Gkaragkounis. On the approximation power of two-layer networks of random ReLUs. *Conference on Learning Theory*, 134:2423–2461, 2021.

- [18] Hailong Hu and Jun Pang. Membership inference of diffusion models. *arXiv preprint arXiv:2301.09956*, 2023.
- [19] Matthew Jagielski, Om Thakkar, Florian Tramer, Daphne Ippolito, Katherine Lee, Nicholas Carlini, Eric Wallace, Shuang Song, Abhradeep Guha Thakurta, Nicolas Papernot, and Chiyuan Zhang. Measuring forgetting of memorized training examples. In *International Conference on Learning Representations*, 2023.
- [20] 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.
- [21] Frederic Koehler, Alexander Heckett, and Andrej Risteski. Statistical efficiency of score matching: The view from isoperimetry. In *International Conference on Learning Representations*, 2023.
- [22] Haoying Li, Yifan Yang, Meng Chang, Shiqi Chen, Huajun Feng, Zhihai Xu, Qi Li, and Yueting Chen. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing*, 479:47–59, 2022.
- [23] Zhong Li, Jiequn Han, Weinan E, and Qianxiao Li. On the curse of memory in recurrent neural networks: Approximation and optimization analysis. In *International Conference on Learning Representations*, 2021.
- [24] Zhong Li, Jiequn Han, Weinan E, and Qianxiao Li. Approximation and optimization theory for linear continuous-time recurrent neural networks. *Journal of Machine Learning Research*, 23(42):1–85, 2022.
- [25] Guan-Horng Liu, Arash Vahdat, De-An Huang, Evangelos Theodorou, Weili Nie, and Anima Anandkumar. I<sup>2</sup>SB: Image-to-image Schrödinger bridge. *International Conference on Machine Learning*, 202:22042–22062, 2023.
- [26] Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint arXiv:2209.14577*, 2022.
- [27] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *International Conference on Learning Representations*, 2023.
- [28] Tomoya Matsumoto, Takayuki Miura, and Naoto Yanai. Membership inference attacks against diffusion models. In *IEEE Security and Privacy Workshops*, pages 77–83, 2023.
- [29] Laurence Illing Midgley, Vincent Stimper, Gregor N. C. Simm, Bernhard Schölkopf, and José Miguel Hernández-Lobato. Flow annealed importance sampling bootstrap. In *International Conference on Learning Representations*, 2023.
- [30] Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.
- [31] Bernt Øksendal. *Stochastic Differential Equations*. Springer, 2003.
- [32] Jakiw Pidstrigach. Score-based generative models detect manifolds. *Advances in Neural Information Processing Systems*, 35:35852–35865, 2022.
- [33] Vadim Popov, Ivan Vovk, Vladimir Gogoryan, Tasnima Sadekova, and Mikhail Kudinov. GradTts: A diffusion probabilistic model for text-to-speech. *International Conference on Machine Learning*, 139:8599–8608, 2021.
- [34] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [35] Kashif Rasul, Calvin Seward, Ingmar Schuster, and Roland Vollgraf. Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. *International Conference on Machine Learning*, 139:8857–8868, 2021.

- [36] 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.
- [37] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015.
- [38] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo-Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022.
- [39] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4713–4726, 2023.
- [40] Simo Särkkä and Arno Solin. *Applied Stochastic Differential Equations*, volume 10. Cambridge University Press, 2019.
- [41] Vikash Sehwal, Caner Hazirbas, Albert Gordo, Firat Ozgenel, and Cristian Canton. Generating high fidelity data from low-density regions using diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11492–11501, 2022.
- [42] Kulin Shah, Sitan Chen, and Adam Klivans. Learning mixtures of gaussians using the DDPM objective. *Advances in Neural Information Processing Systems*, 36, 2023.
- [43] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- [44] Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion Schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36, 2023.
- [45] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. *International Conference on Machine Learning*, 37:2256–2265, 2015.
- [46] 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.
- [47] Vignesh Ram Somnath, Matteo Pariset, Ya-Ping Hsieh, Maria Rodriguez Martinez, Andreas Krause, and Charlotte Bunne. Aligned diffusion Schrödinger bridges. *Conference on Uncertainty in Artificial Intelligence*, 216:1985–1995, 2023.
- [48] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *International Conference on Machine Learning*, 202:32211–32252, 2023.
- [49] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems*, 34:1415–1428, 2021.
- [50] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019.
- [51] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in Neural Information Processing Systems*, 33:12438–12448, 2020.
- [52] Yang Song, Sahaj Garg, Jiabin Shi, and Stefano Ermon. Sliced score matching: A scalable approach to density and score estimation. *Conference on Uncertainty in Artificial Intelligence*, 115:574–584, 2020.

- [53] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *International Conference on Learning Representations*, 2022.
- [54] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- [55] Ramon Van Handel. Probability in high dimension. *Lecture Notes (Princeton University)*, 2014.
- [56] Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving Schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.
- [57] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.
- [58] Nikhil Vyas, Sham M. Kakade, and Boaz Barak. On provable copyright protection for generative models. *International Conference on Machine Learning*, 202:35277–35299, 2023.
- [59] Gefei Wang, Yuling Jiao, Qian Xu, Yang Wang, and Can Yang. Deep generative learning via Schrödinger bridge. *International Conference on Machine Learning*, 139:10794–10804, 2021.
- [60] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7623–7633, 2023.
- [61] Lei Wu and Weijie J. Su. The implicit regularization of dynamical stability in stochastic gradient descent. *International Conference on Machine Learning*, 202:37656–37684, 2023.
- [62] Yixin Wu, Ning Yu, Zheng Li, Michael Backes, and Yang Zhang. Membership inference attacks against text-to-image generation models. *arXiv preprint arXiv:2210.00968*, 2022.
- [63] Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi S. Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. In *International Conference on Learning Representations*, 2022.
- [64] Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. In *International Conference on Learning Representations*, 2022.
- [65] Hongkang Yang. A mathematical framework for learning probability distributions. *Journal of Machine Learning*, 1(4):373–431, 2022.
- [66] Hongkang Yang and Weinan E. Generalization and memorization: The bias potential model. *Mathematical and Scientific Machine Learning*, 145:1013–1043, 2022.
- [67] Hongkang Yang and Weinan E. Generalization error of GAN from the discriminator’s perspective. *Research in the Mathematical Sciences*, 9(1):1–31, 2022.
- [68] Ruihan Yang, Prakhar Srivastava, and Stephan Mandt. Diffusion probabilistic modeling for video generation. *Entropy*, 25(10):1469, 2023.
- [69] Jongmin Yoon, Sung Ju Hwang, and Juho Lee. Adversarial purification with score-based generative models. *International Conference on Machine Learning*, 139:12062–12072, 2021.
- [70] Chiyuan Zhang, Daphne Ippolito, Katherine Lee, Matthew Jagielski, Florian Tramèr, and Nicholas Carlini. Counterfactual memorization in neural language models. *Advances in Neural Information Processing Systems*, 36, 2023.

## A Technical Results and Proofs

### A.1 Data-Independent Generalization Gap

To derive the theorem for the generalization error of this score-based generative model, we first give the following lemmas.

**Lemma 1** (Forward perturbation estimates). *Consider the forward diffusion process with the linear drift coefficient (2). For any  $\delta > 0$ ,  $\delta \ll 1$ , with the probability of at least  $1 - \delta$ , we have*

$$\|\mathbf{x}(t)\|_\infty \lesssim C_T \left( \|\mathbf{x}(0)\|_\infty + \sqrt{\log(1/\delta^2)} \right), \quad (18)$$

where  $C_T := \max_{t \in [0, T]} \{r(t), r(t)v(t)\}$ .

*Proof.* When the drift coefficient  $\mathbf{f}(\cdot, t) : \mathbb{R}^d \mapsto \mathbb{R}^d$  is linear to  $\mathbf{x}$ , i.e.,  $\mathbf{f}(\mathbf{x}, t) = f(t)\mathbf{x}$ , the transition kernel  $p_{t|0}$  has a closed form (8)

$$p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0)) = \mathcal{N}(\mathbf{x}(t); r(t)\mathbf{x}(0), r^2(t)v^2(t)\mathbf{I}_d) \quad (19)$$

with  $r(t) := e^{\int_0^t f(\zeta)d\zeta}$ ,  $v(t) := \sqrt{\int_0^t \frac{g^2(\zeta)}{r^2(\zeta)} d\zeta}$ . Hence, we have

$$\mathbf{x}(t) = r(t)\mathbf{x}(0) + r(t)v(t)\mathbf{z}, \quad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d). \quad (20)$$

For any  $\epsilon \sim \mathcal{N}(0, 1)$ ,  $c > 1$ , we have

$$\begin{aligned} \mathbb{P}\{\epsilon : |\epsilon| > c\} &= 2 \int_c^{+\infty} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \\ &\leq \frac{1}{\sqrt{2\pi}} \int_c^{+\infty} 2xe^{-x^2/2} dx = \frac{1}{\sqrt{2\pi}} \int_{c^2}^{+\infty} e^{-x/2} dx = \sqrt{\frac{2}{\pi}} e^{-c^2/2}. \end{aligned}$$

Let  $\delta = \Theta(e^{-c^2/2})$ , we get

$$\mathbb{P}\{\epsilon : |\epsilon| \leq \sqrt{\log(1/\delta^2)}\} \geq 1 - \delta. \quad (21)$$

Hence, for any  $\delta \in (0, 1)$  with  $\delta \ll 1$ , with the probability of at least  $1 - \delta$ , we have

$$\|\mathbf{x}(t)\|_\infty \lesssim C_T \left( \|\mathbf{x}(0)\|_\infty + \sqrt{\log(1/\delta^2)} \right) \quad (22)$$

with  $C_T := \max_{t \in [0, T]} \{r(t), r(t)v(t)\}$ . The proof is completed.  $\square$

**Lemma 2** (Theorem 1 in [49]). *We have*

$$D_{\text{KL}}(p_0 \| p_{0, \hat{\theta}_n(\tau)}) \leq \tilde{\mathcal{L}}(\hat{\theta}_n(\tau); g^2(\cdot)) + D_{\text{KL}}(p_T \| \pi).$$

**Lemma 3.** *Both of the loss objectives  $\tilde{\mathcal{L}}(\boldsymbol{\theta}; \lambda(\cdot))$  and  $\tilde{\mathcal{L}}(\bar{\boldsymbol{\theta}}; \lambda(\cdot))$  are quadratic (and hence convex).*

*Proof.* The convexity arises from the following points: (i) The score matching loss objectives  $\tilde{\mathcal{L}}(\boldsymbol{\theta}; \lambda(\cdot))$  (defined in (7)) and  $\tilde{\mathcal{L}}(\bar{\boldsymbol{\theta}}; \lambda(\cdot))$  are  $L^2$ -metrics between the score network model and target score function; (ii) The score networks are defined as random feature models (see (12) and (13)) that are linear to trainable parameters. Therefore, using some trace techniques and basic variational calculations, it is not hard to derive the fact that the loss objectives are quadratic and hence convex with respect to trainable parameters.

(1) For  $\tilde{\mathcal{L}}(\boldsymbol{\theta}; \lambda(\cdot))$ , recall that  $\mathbf{s}_{t, \boldsymbol{\theta}}(\mathbf{x}(t)) = \frac{1}{m} \mathbf{A} \sigma(\mathbf{W} \mathbf{x}(t) + \mathbf{U} \mathbf{e}(t))$ , and let  $\mathbf{s}_t(\mathbf{x}(t)) := \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t))$ ,  $\mathbf{h}_1(\mathbf{x}, t) := (\sqrt{\lambda(t)}/\sqrt{m}) \sigma(\mathbf{W} \mathbf{x} + \mathbf{U} \mathbf{e}(t))$ ,  $\mathbf{h}_2(\mathbf{x}, t) := \sqrt{\lambda(t)} \mathbf{s}_t(\mathbf{x})$ , we have

$$\begin{aligned} \tilde{\mathcal{L}}(\boldsymbol{\theta}; \lambda(\cdot)) &= \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \mathbf{h}_1^\top(\mathbf{x}(t), t) (\mathbf{A}/\sqrt{m})^\top (\mathbf{A}/\sqrt{m}) \mathbf{h}_1(\mathbf{x}(t), t) \right. \\ &\quad \left. - 2\mathbf{h}_2^\top(\mathbf{x}(t), t) (\mathbf{A}/\sqrt{m}) \mathbf{h}_1(\mathbf{x}(t), t) \right] + \text{constant}. \end{aligned}$$

Since for any  $\mathbf{h}, \bar{\mathbf{h}}, \mathbf{B}$ , we have

$$\begin{aligned}\mathbb{E}_t \mathbb{E}_{\mathbf{x}(t)} [\mathbf{h}^\top(\mathbf{x}(t), t) \mathbf{B} \bar{\mathbf{h}}(\mathbf{x}(t), t)] &= \mathbb{E}_t \mathbb{E}_{\mathbf{x}(t)} [\text{trace}(\mathbf{B} \bar{\mathbf{h}}(\mathbf{x}(t), t) \mathbf{h}^\top(\mathbf{x}(t), t))] \\ &= \text{trace}(\mathbf{B} \mathbb{E}_t \mathbb{E}_{\mathbf{x}(t)} [\bar{\mathbf{h}}(\mathbf{x}(t), t) \mathbf{h}^\top(\mathbf{x}(t), t)]),\end{aligned}$$

we further get

$$\tilde{\mathcal{L}}(\boldsymbol{\theta}; \lambda(\cdot)) = \frac{1}{m} \text{trace}(\mathbf{A}^\top \mathbf{A} \mathbf{B}_1) - \frac{2}{\sqrt{m}} \text{trace}(\mathbf{A} \mathbf{B}_2) + \text{constant},$$

where

$$\mathbf{B}_1 := \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} [\mathbf{h}_1(\mathbf{x}(t), t) \mathbf{h}_1^\top(\mathbf{x}(t), t)], \quad \mathbf{B}_2 := \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} [\mathbf{h}_1(\mathbf{x}(t), t) \mathbf{h}_2^\top(\mathbf{x}(t), t)].$$

Here,  $\mathbf{B}_1$  is a positive semi-definite matrix, since  $\mathbf{v}^\top \mathbf{B}_1 \mathbf{v} = \mathbb{E}_t \mathbb{E}_{\mathbf{x}(t)} [(\mathbf{v}^\top \mathbf{h}_1(\mathbf{x}(t), t))^2] \geq 0$  for any  $\mathbf{v}$ . Notice that for any  $\mathbf{A}, \mathbf{B}$ ,

$$\begin{aligned}\text{trace}(\mathbf{A}^\top \mathbf{A} \mathbf{B}) &= \text{trace}(\mathbf{A} \mathbf{B} \mathbf{A}^\top) = \sum_{i, j} \mathbf{B}_{ij} (\mathbf{A}_{:,j})^\top \mathbf{A}_{:,i} = \text{vec}(\mathbf{A})^\top (\mathbf{B} \otimes \mathbf{I}) \text{vec}(\mathbf{A}), \\ \text{trace}(\mathbf{A} \mathbf{B}) &= \sum_j (\mathbf{A}_{:,j})^\top (\mathbf{B}^\top)_{:,j} = \text{vec}(\mathbf{A})^\top \text{vec}(\mathbf{B}^\top),\end{aligned}$$

where  $\otimes$  denotes the Kronecker product. Hence

$$\tilde{\mathcal{L}}(\boldsymbol{\theta}; \lambda(\cdot)) = \frac{1}{m} \text{vec}(\mathbf{A})^\top (\mathbf{B}_1 \otimes \mathbf{I}) \text{vec}(\mathbf{A}) - \frac{2}{\sqrt{m}} \text{vec}(\mathbf{B}_2^\top)^\top \text{vec}(\mathbf{A}) + \text{constant} \quad (23)$$

is a quadratic function. It is straightforward to show that the eigenvalues of  $\mathbf{B}_1 \otimes \mathbf{I}$  are the same as  $\mathbf{B}_1$  but with multiplicity,<sup>7</sup> implying that  $\mathbf{B}_1 \otimes \mathbf{I}$  is also positive semi-definite. Therefore,

$$\nabla_{\boldsymbol{\theta}}^2 \tilde{\mathcal{L}}(\boldsymbol{\theta}; \lambda(\cdot)) = \nabla_{\text{vec}(\mathbf{A})}^2 \tilde{\mathcal{L}}(\boldsymbol{\theta}; \lambda(\cdot)) = 2(\mathbf{B}_1 \otimes \mathbf{I})$$

is positive semi-definite, i.e., the loss is convex with respect to trainable parameters.

(2) For  $\bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}; \lambda(\cdot))$ , notice that

$$\begin{aligned}\|\bar{\mathbf{s}}_{t, \bar{\boldsymbol{\theta}}}(\mathbf{x})\|_2^2 &= \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [\sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t)) \mathbf{a}^\top(\mathbf{w}, \mathbf{u})] \mathbb{E}_{(\mathbf{w}', \mathbf{u}') \sim \rho_0} [\mathbf{a}(\mathbf{w}', \mathbf{u}') \sigma(\mathbf{w}'^\top \mathbf{x} + \mathbf{u}'^\top \mathbf{e}(t))] \\ &= \mathbb{E}_{(\mathbf{w}, \mathbf{u}), (\mathbf{w}', \mathbf{u}') \sim \rho_0} [\mathbf{a}^\top(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t)) \sigma(\mathbf{w}'^\top \mathbf{x} + \mathbf{u}'^\top \mathbf{e}(t)) \mathbf{a}(\mathbf{w}', \mathbf{u}')].\end{aligned}$$

Let  $\mathbf{v} := (\mathbf{w}, \mathbf{u})$ ,  $\mathbf{v}' := (\mathbf{w}', \mathbf{u}')$ ,  $\mathbf{z}(t) := (\mathbf{x}^\top(t), \mathbf{e}^\top(t))^\top$ , we get

$$\begin{aligned}\bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}; \lambda(\cdot)) &= \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} [\lambda(t) (\|\bar{\mathbf{s}}_{t, \bar{\boldsymbol{\theta}}}(\mathbf{x}(t))\|_2^2 - 2\mathbf{s}_t^\top(\mathbf{x}(t)) \bar{\mathbf{s}}_{t, \bar{\boldsymbol{\theta}}}(\mathbf{x}(t)) + \|\mathbf{s}_t(\mathbf{x}(t))\|_2^2)] \\ &= \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} [\lambda(t) (\mathbb{E}_{\mathbf{v}, \mathbf{v}' \sim \rho_0} [\mathbf{a}^\top(\mathbf{v}) \sigma(\mathbf{v}^\top \mathbf{z}(t)) \sigma(\mathbf{v}'^\top \mathbf{z}(t)) \mathbf{a}(\mathbf{v}')] \\ &\quad - 2\mathbf{s}_t^\top(\mathbf{x}(t)) \mathbb{E}_{\mathbf{v} \sim \rho_0} [\mathbf{a}(\mathbf{v}) \sigma(\mathbf{v}^\top \mathbf{z}(t))])] + \text{constant} \\ &= \mathbb{E}_{\mathbf{v}, \mathbf{v}' \sim \rho_0} [\mathbf{a}^\top(\mathbf{v}) (\mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} [\lambda(t) \sigma(\mathbf{v}^\top \mathbf{z}(t)) \sigma(\mathbf{v}'^\top \mathbf{z}(t))]) \mathbf{a}(\mathbf{v}')] \\ &\quad - 2\mathbb{E}_{\mathbf{v} \sim \rho_0} [(\mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} [\lambda(t) \mathbf{s}_t^\top(\mathbf{x}(t)) \sigma(\mathbf{v}^\top \mathbf{z}(t))]) \mathbf{a}(\mathbf{v})] + \text{constant}.\end{aligned}$$

<sup>7</sup>In fact, if  $\mathbf{A} \in \mathbb{R}^{n \times n}$ ,  $\mathbf{B} \in \mathbb{R}^{m \times m}$  have the eigenvalues  $\{\nu_i\}_{i=1}^n$ ,  $\{\mu_j\}_{j=1}^m$ , respectively, the eigenvalues of  $\mathbf{A} \otimes \mathbf{B}$  are  $\nu_i \mu_j$ ,  $i = 1, \dots, n$ ,  $j = 1, \dots, m$ . This is due to the following argument. By Jordan–Chevalley decomposition, there exist invertible matrices  $\mathbf{P}, \mathbf{Q}$  such that  $\mathbf{A} = \mathbf{P} \boldsymbol{\Lambda} \mathbf{P}^{-1}$ ,  $\mathbf{B} = \mathbf{Q} \boldsymbol{\Delta} \mathbf{Q}^{-1}$ , where  $\boldsymbol{\Lambda}, \boldsymbol{\Delta}$  are upper triangular matrices. Therefore, we have

$$\mathbf{A} \otimes \mathbf{B} = (\mathbf{P} \boldsymbol{\Lambda} \mathbf{P}^{-1}) \otimes (\mathbf{Q} \boldsymbol{\Delta} \mathbf{Q}^{-1}) = (\mathbf{P} \otimes \mathbf{Q}) (\boldsymbol{\Lambda} \otimes \boldsymbol{\Delta}) (\mathbf{P}^{-1} \otimes \mathbf{Q}^{-1}) = (\mathbf{P} \otimes \mathbf{Q}) (\boldsymbol{\Lambda} \otimes \boldsymbol{\Delta}) (\mathbf{P} \otimes \mathbf{Q})^{-1}.$$

That is, the matrices  $\mathbf{A} \otimes \mathbf{B}$  and  $\boldsymbol{\Lambda} \otimes \boldsymbol{\Delta}$  are similar. Notice that  $\boldsymbol{\Lambda} \otimes \boldsymbol{\Delta}$  is still an upper triangular matrix with diagonal elements  $\nu_i \mu_j$ ,  $i = 1, \dots, n$ ,  $j = 1, \dots, m$ , we obtain the desired result.

Then for any  $\phi \in L^2(\rho_0)$ , by symmetry we have

$$\begin{aligned} \delta \bar{\mathcal{L}}(\cdot; \lambda(\cdot))[\bar{\theta}, \phi] &= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left( \bar{\mathcal{L}}(\bar{\theta} + \epsilon \phi; \lambda(\cdot)) - \bar{\mathcal{L}}(\bar{\theta}; \lambda(\cdot)) \right) \\ &= \mathbb{E}_{\mathbf{v}, \mathbf{v}' \sim \rho_0} \left[ \phi^\top(\mathbf{v}) \left( \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \sigma(\mathbf{v}^\top \mathbf{z}(t)) \sigma(\mathbf{v}'^\top \mathbf{z}(t)) \right] \right) \mathbf{a}(\mathbf{v}') \right. \\ &\quad \left. + \mathbf{a}^\top(\mathbf{v}) \left( \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \sigma(\mathbf{v}^\top \mathbf{z}(t)) \sigma(\mathbf{v}'^\top \mathbf{z}(t)) \right] \right) \phi(\mathbf{v}') \right] \\ &\quad - 2 \mathbb{E}_{\mathbf{v} \sim \rho_0} \left[ \left( \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \mathbf{s}_t^\top(\mathbf{x}(t)) \sigma(\mathbf{v}^\top \mathbf{z}(t)) \right] \right) \phi(\mathbf{v}) \right] \\ &= 2 \mathbb{E}_{\mathbf{v}, \mathbf{v}' \sim \rho_0} \left[ \mathbf{a}^\top(\mathbf{v}') \left( \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \sigma(\mathbf{v}^\top \mathbf{z}(t)) \sigma(\mathbf{v}'^\top \mathbf{z}(t)) \right] \right) \phi(\mathbf{v}) \right] \\ &\quad - 2 \mathbb{E}_{\mathbf{v} \sim \rho_0} \left[ \left( \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \mathbf{s}_t^\top(\mathbf{x}(t)) \sigma(\mathbf{v}^\top \mathbf{z}(t)) \right] \right) \phi(\mathbf{v}) \right] \\ &= 2 \left\langle \mathbb{E}_{\mathbf{v}' \sim \rho_0} \left[ K(\mathbf{v}, \mathbf{v}'; \lambda(\cdot)) \mathbf{a}(\mathbf{v}') \right] - \mathbf{k}(\mathbf{v}; \lambda(\cdot)), \phi(\mathbf{v}) \right\rangle_{L^2(\rho_0)}, \end{aligned}$$

where

$$\begin{aligned} K(\mathbf{v}, \mathbf{v}'; \lambda(\cdot)) &:= \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \sigma(\mathbf{v}^\top \mathbf{z}(t)) \sigma(\mathbf{v}'^\top \mathbf{z}(t)) \right], \\ \mathbf{k}(\mathbf{v}; \lambda(\cdot)) &:= \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \mathbf{s}_t(\mathbf{x}(t)) \sigma(\mathbf{v}^\top \mathbf{z}(t)) \right]. \end{aligned}$$

This yields

$$\frac{\delta \bar{\mathcal{L}}(\cdot; \lambda(\cdot))}{\delta \bar{\theta}} = 2 \mathbb{E}_{\mathbf{v}' \sim \rho_0} \left[ K(\mathbf{v}, \mathbf{v}'; \lambda(\cdot)) \mathbf{a}(\mathbf{v}') \right] - 2 \mathbf{k}(\mathbf{v}; \lambda(\cdot)), \quad (24)$$

and

$$\begin{aligned} &\bar{\mathcal{L}}(\bar{\theta}_1; \lambda(\cdot)) + \left\langle \frac{\delta \bar{\mathcal{L}}(\cdot; \lambda(\cdot))}{\delta \bar{\theta}} \Bigg|_{\bar{\theta} = \bar{\theta}_1}, \bar{\theta}_2 - \bar{\theta}_1 \right\rangle_{L^2(\rho_0)} - \bar{\mathcal{L}}(\bar{\theta}_2; \lambda(\cdot)) \\ &= \mathbb{E}_{\mathbf{v}, \mathbf{v}' \sim \rho_0} \left[ \mathbf{a}_1^\top(\mathbf{v}) K(\mathbf{v}, \mathbf{v}'; \lambda(\cdot)) \mathbf{a}_1(\mathbf{v}') \right] - 2 \mathbb{E}_{\mathbf{v} \sim \rho_0} \left[ \mathbf{k}^\top(\mathbf{v}; \lambda(\cdot)) \mathbf{a}_1(\mathbf{v}) \right] \\ &\quad + 2 \left\langle \mathbb{E}_{\mathbf{v}' \sim \rho_0} \left[ K(\mathbf{v}, \mathbf{v}'; \lambda(\cdot)) \mathbf{a}_1(\mathbf{v}') \right] - \mathbf{k}(\mathbf{v}; \lambda(\cdot)), \mathbf{a}_2(\mathbf{v}) - \mathbf{a}_1(\mathbf{v}) \right\rangle_{L^2(\rho_0)} \\ &\quad - \mathbb{E}_{\mathbf{v}, \mathbf{v}' \sim \rho_0} \left[ \mathbf{a}_2^\top(\mathbf{v}) K(\mathbf{v}, \mathbf{v}'; \lambda(\cdot)) \mathbf{a}_2(\mathbf{v}') \right] + 2 \mathbb{E}_{\mathbf{v} \sim \rho_0} \left[ \mathbf{k}^\top(\mathbf{v}; \lambda(\cdot)) \mathbf{a}_2(\mathbf{v}) \right] \\ &= - \mathbb{E}_{\mathbf{v}, \mathbf{v}' \sim \rho_0} \left[ (\mathbf{a}_2(\mathbf{v}) - \mathbf{a}_1(\mathbf{v}))^\top K(\mathbf{v}, \mathbf{v}'; \lambda(\cdot)) (\mathbf{a}_2(\mathbf{v}') - \mathbf{a}_1(\mathbf{v}')) \right] \\ &= - \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \mathbb{E}_{\mathbf{v}, \mathbf{v}' \sim \rho_0} \left[ (\mathbf{a}_2(\mathbf{v}) - \mathbf{a}_1(\mathbf{v}))^\top \sigma(\mathbf{v}^\top \mathbf{z}(t)) \sigma(\mathbf{v}'^\top \mathbf{z}(t)) (\mathbf{a}_2(\mathbf{v}') - \mathbf{a}_1(\mathbf{v}')) \right] \right] \\ &= - \mathbb{E}_{t \sim \mathcal{U}(0, T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \left\| \mathbb{E}_{\mathbf{v} \sim \rho_0} \left[ (\mathbf{a}_2(\mathbf{v}) - \mathbf{a}_1(\mathbf{v})) \sigma(\mathbf{v}^\top \mathbf{z}(t)) \right] \right\|_2^2 \right] \leq 0, \end{aligned}$$

hence the functional  $\bar{\mathcal{L}}(\bar{\theta}; \lambda(\cdot))$  is convex with respect to  $\bar{\theta}$  (given any positive weighting function  $\lambda(\cdot)$ ). The proof is completed.  $\square$

Since the weighting function is fixed in our analysis, we omit the notation  $\lambda(\cdot)$  without ambiguity in the following contents.

**Lemma 4.** For any  $\tau > 0$  and  $\theta, \bar{\theta}$ , we have

$$\bar{\mathcal{L}}(\bar{\theta}(\tau)) - \bar{\mathcal{L}}(\bar{\theta}) \lesssim \frac{\|\bar{\mathbf{s}}_{0, \bar{\theta}_0}\|_{\mathcal{H}}^2 + \|\bar{\mathbf{s}}_{0, \bar{\theta}}\|_{\mathcal{H}}^2}{\tau}, \quad \tilde{\mathcal{L}}(\theta(\tau)) - \tilde{\mathcal{L}}(\theta) \lesssim \frac{\|\mathbf{s}_{0, \theta_0}\|_{\mathcal{H}}^2 + \|\mathbf{s}_{0, \theta}\|_{\mathcal{H}}^2}{\tau}.$$

*Proof.* (1) For the loss objective  $\bar{\mathcal{L}}$ , we define the Lyapunov function

$$\bar{E}(\tau) := \tau \left( \bar{\mathcal{L}}(\bar{\theta}(\tau)) - \bar{\mathcal{L}}(\bar{\theta}) \right) + \frac{1}{2} \|\mathbf{a}_\tau - \mathbf{a}\|_{L^2(\rho_0)}^2.$$

Then, we have

$$\begin{aligned} \frac{d}{d\tau} \bar{E}(\tau) &= \left( \bar{\mathcal{L}}(\bar{\theta}(\tau)) - \bar{\mathcal{L}}(\bar{\theta}) \right) + \tau \cdot \frac{d}{d\tau} \bar{\mathcal{L}}(\bar{\theta}(\tau)) + \left\langle \mathbf{a}_\tau - \mathbf{a}, \frac{d}{d\tau} \mathbf{a}_\tau \right\rangle_{L^2(\rho_0)} \\ &\leq \left( \bar{\mathcal{L}}(\bar{\theta}(\tau)) - \bar{\mathcal{L}}(\bar{\theta}) \right) - \left\langle \mathbf{a}_\tau - \mathbf{a}, \frac{\delta \bar{\mathcal{L}}}{\delta \bar{\theta}} \Bigg|_{\bar{\theta} = \bar{\theta}(\tau)} \right\rangle_{L^2(\rho_0)}, \end{aligned}$$

where the last inequality holds since

$$\frac{d}{d\tau} \bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}(\tau)) = \left\langle \frac{\delta \bar{\mathcal{L}}}{\delta \bar{\boldsymbol{\theta}}} \Big|_{\bar{\boldsymbol{\theta}}=\bar{\boldsymbol{\theta}}(\tau)}, \frac{d}{d\tau} \bar{\boldsymbol{\theta}}(\tau) \right\rangle_{L^2(\rho_0)} = - \left\langle \frac{\delta \bar{\mathcal{L}}}{\delta \bar{\boldsymbol{\theta}}} \Big|_{\bar{\boldsymbol{\theta}}=\bar{\boldsymbol{\theta}}(\tau)}, \frac{\delta \bar{\mathcal{L}}}{\delta \bar{\boldsymbol{\theta}}} \Big|_{\bar{\boldsymbol{\theta}}=\bar{\boldsymbol{\theta}}(\tau)} \right\rangle_{L^2(\rho_0)} \leq 0.$$

By convexity, for any  $\tau_1, \tau_2$ , it holds that

$$\bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}(\tau_1)) + \left\langle \mathbf{a}_{\tau_2} - \mathbf{a}_{\tau_1}, \frac{\delta \bar{\mathcal{L}}}{\delta \bar{\boldsymbol{\theta}}} \Big|_{\bar{\boldsymbol{\theta}}=\bar{\boldsymbol{\theta}}(\tau_1)} \right\rangle_{L^2(\rho_0)} \leq \bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}(\tau_2)),$$

hence  $\frac{d}{d\tau} \bar{E}(\tau) \leq 0$ . We conclude that  $\bar{E}(\tau) \leq \bar{E}(0)$ , or equivalently

$$\tau \left( \bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}(\tau)) - \bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}) \right) + \frac{1}{2} \|\mathbf{a}_\tau - \mathbf{a}\|_{L^2(\rho_0)}^2 \leq \frac{1}{2} \|\mathbf{a}_0 - \mathbf{a}\|_{L^2(\rho_0)}^2.$$

Therefore, note that  $\|\bar{\mathbf{s}}_{0,\boldsymbol{\theta}}\|_{\mathcal{H}}^2 := \mathbb{E}_{\rho_0}[\|\mathbf{a}\|_2^2]$  (let  $\mathcal{H} := \mathcal{H}_{k,\rho_0}$ ), we obtain

$$\bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}(\tau)) - \bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}) \lesssim \frac{\|\mathbf{a}_0 - \mathbf{a}\|_{L^2(\rho_0)}^2}{\tau} \lesssim \frac{\|\bar{\mathbf{s}}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}}^2 + \|\bar{\mathbf{s}}_{0,\boldsymbol{\theta}}\|_{\mathcal{H}}^2}{\tau},$$

which gives the desired estimate.

(2) For the loss objective  $\tilde{\mathcal{L}}$ , the argument is almost the same, except replacing the Lyapunov function by

$$E(\tau) := \tau \left( \tilde{\mathcal{L}}(\boldsymbol{\theta}(\tau)) - \tilde{\mathcal{L}}(\boldsymbol{\theta}) \right) + \frac{1}{2m} \|\mathbf{A}_\tau - \mathbf{A}\|_F^2,$$

and  $\langle \cdot, \cdot \rangle_{L^2(\rho_0)}$  by  $\langle \cdot, \cdot \rangle$ . Note that  $\boldsymbol{\theta} = \text{vec}(\mathbf{A})/\sqrt{m}$  and  $\|\mathbf{s}_{0,\boldsymbol{\theta}}\|_{\mathcal{H}}^2 = \|\mathbf{A}\|_F^2/m$ , we obtain

$$\tilde{\mathcal{L}}(\boldsymbol{\theta}(\tau)) - \tilde{\mathcal{L}}(\boldsymbol{\theta}) \lesssim \frac{\|\mathbf{A}_0 - \mathbf{A}\|_F^2}{m\tau} \lesssim \frac{\|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}}^2 + \|\mathbf{s}_{0,\boldsymbol{\theta}}\|_{\mathcal{H}}^2}{\tau},$$

which completes the proof.  $\square$

**Lemma 5.** Suppose that the loss objectives  $\tilde{\mathcal{L}}, \tilde{\mathcal{L}}^{(n)}, \bar{\mathcal{L}}, \bar{\mathcal{L}}^{(n)}$  are bounded at the initialization, then for any  $\tau > 0$ , we have

$$\|\mathbf{s}_{0,\boldsymbol{\theta}(\tau)}\|_{\mathcal{H}}, \|\mathbf{s}_{0,\hat{\boldsymbol{\theta}}_n(\tau)}\|_{\mathcal{H}} \lesssim \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\frac{\tau}{m}}, \quad \|\bar{\mathbf{s}}_{0,\bar{\boldsymbol{\theta}}(\tau)}\|_{\mathcal{H}}, \|\bar{\mathbf{s}}_{0,\hat{\bar{\boldsymbol{\theta}}}_n(\tau)}\|_{\mathcal{H}} \lesssim \|\bar{\mathbf{s}}_{0,\bar{\boldsymbol{\theta}}_0}\|_{\mathcal{H}} + \sqrt{\tau}.$$

*Proof.* It's sufficient to prove  $\|\bar{\mathbf{s}}_{0,\bar{\boldsymbol{\theta}}(\tau)}\|_{\mathcal{H}} \lesssim \sqrt{\tau}$ , and the rest part follows similarly.

Since

$$\|\mathbf{a}_\tau\|_{L^2(\rho_0)} \frac{d}{d\tau} \|\mathbf{a}_\tau\|_{L^2(\rho_0)} = \frac{1}{2} \frac{d}{d\tau} \|\mathbf{a}_\tau\|_{L^2(\rho_0)}^2 = \left\langle \mathbf{a}_\tau, \frac{d}{d\tau} \mathbf{a}_\tau \right\rangle_{L^2(\rho_0)} = \left\langle \mathbf{a}_\tau, -\frac{\delta \bar{\mathcal{L}}}{\delta \bar{\boldsymbol{\theta}}} \Big|_{\bar{\boldsymbol{\theta}}=\bar{\boldsymbol{\theta}}(\tau)} \right\rangle_{L^2(\rho_0)},$$

applying Cauchy–Schwartz inequality yields

$$\begin{aligned} \frac{d}{d\tau} \|\mathbf{a}_\tau\|_{L^2(\rho_0)} &= \left\langle \frac{\mathbf{a}_\tau}{\|\mathbf{a}_\tau\|_{L^2(\rho_0)}}, -\frac{\delta \bar{\mathcal{L}}}{\delta \bar{\boldsymbol{\theta}}} \Big|_{\bar{\boldsymbol{\theta}}=\bar{\boldsymbol{\theta}}(\tau)} \right\rangle_{L^2(\rho_0)} \\ &\leq \left\| \frac{\delta \bar{\mathcal{L}}}{\delta \bar{\boldsymbol{\theta}}} \Big|_{\bar{\boldsymbol{\theta}}=\bar{\boldsymbol{\theta}}(\tau)} \right\|_{L^2(\rho_0)} = \sqrt{-\frac{d}{d\tau} \bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}(\tau))}. \end{aligned}$$

Thus, again by Cauchy–Schwartz inequality, for any  $\tau > \tau_0 \geq 0$ , we have

$$\begin{aligned} \|\mathbf{a}_\tau\|_{L^2(\rho_0)} - \|\mathbf{a}_{\tau_0}\|_{L^2(\rho_0)} &\leq \int_{\tau_0}^{\tau} \sqrt{-\frac{d}{ds} \bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}(s))} ds \\ &\leq \sqrt{\tau - \tau_0} \sqrt{-\bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}(\tau)) + \bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}(\tau_0))} \\ &\leq \sqrt{\tau} \sqrt{\bar{\mathcal{L}}(\bar{\boldsymbol{\theta}}(0))}. \end{aligned}$$

By choosing  $\tau_0 = 0$ , we have  $\|\mathbf{a}_\tau\|_{L^2(\rho_0)} \lesssim \|\bar{\mathbf{s}}_{0,\bar{\theta}(0)}\|_{\mathcal{H}} + \sqrt{\tau}$ , hence  $\|\bar{\mathbf{s}}_{0,\bar{\theta}(\tau)}\|_{\mathcal{H}} \lesssim \|\bar{\mathbf{s}}_{0,\bar{\theta}(0)}\|_{\mathcal{H}} + \sqrt{\tau}$ , which completes the proof.  $\square$

**Lemma 6** (Monte Carlo estimates). *Define the Monte Carlo error*

$$\text{Err}_{\text{MC}} = \text{Err}_{\text{MC}}(\boldsymbol{\theta}, \bar{\boldsymbol{\theta}}; T, \lambda(\cdot)) := \mathbb{E}_{t \sim \mathcal{U}(0,T)} \left[ \lambda(t) \cdot \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \|\mathbf{s}_{t,\boldsymbol{\theta}}(\mathbf{x}(t)) - \bar{\mathbf{s}}_{t,\bar{\boldsymbol{\theta}}}(\mathbf{x}(t))\|_2^2 \right] \right]. \quad (25)$$

Suppose that  $\|\mathbf{x}(0)\|_\infty \leq 1$ , and the trainable parameter  $\mathbf{a}$  and embedding function  $\mathbf{e}(\cdot)$  are both bounded. Then, given any  $\bar{\boldsymbol{\theta}}$ , for any  $\delta > 0$ ,  $\delta \ll 1$ , with the probability of at least  $1 - \delta$ , there exists  $\boldsymbol{\theta}$  such that

$$\text{Err}_{\text{MC}} \lesssim \frac{\log^2(1/\delta^2)}{m} d, \quad (26)$$

where  $\lesssim$  hides universal positive constants only depending on  $T$ .

*Proof.* Fix any  $\bar{\boldsymbol{\theta}}$ . According to Lemma 1, for any  $\delta > 0$ ,  $\delta \ll 1$ , with the probability of at least  $1 - \delta$ , we have

$$\|\mathbf{x}(t)\|_\infty \lesssim C_T \left( 1 + \sqrt{\log(1/\delta^2)} \right) \triangleq C_{T,\delta}. \quad (27)$$

Based on the representation (13), for any  $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_m)^\top \in \mathbb{R}^{m \times d}$ ,  $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_m)^\top \in \mathbb{R}^{m \times d_e}$  with  $(\mathbf{w}_i, \mathbf{u}_i) \sim \rho_0$ ,  $i = 1, \dots, m$ , let  $\mathbf{A} := (\mathbf{a}_1, \dots, \mathbf{a}_m) \in \mathbb{R}^{d \times m}$  with  $\mathbf{a}_i := \mathbf{a}(\mathbf{w}_i, \mathbf{u}_i)$  for  $i = 1, \dots, m$ , and

$$\mathbf{s}_{t,\boldsymbol{\theta}}(\mathbf{x}) := \frac{1}{m} \mathbf{A} \sigma(\mathbf{W} \mathbf{x} + \mathbf{U} \mathbf{e}(t)) = \frac{1}{m} \sum_{i=1}^m \mathbf{a}_i \sigma(\mathbf{w}_i^\top \mathbf{x} + \mathbf{u}_i^\top \mathbf{e}(t)), \quad (28)$$

then  $\mathbb{E}_{\mathbf{W}, \mathbf{U}} [\mathbf{s}_{t,\boldsymbol{\theta}}(\mathbf{x})] = \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [\mathbf{a}(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t))] = \bar{\mathbf{s}}_{t,\bar{\boldsymbol{\theta}}}(\mathbf{x})$ . For  $k = 1, \dots, d$ , let

$$\begin{aligned} Z_{t,k}(\mathbf{W}, \mathbf{U}) &:= \left\| \mathbf{s}_{t,\boldsymbol{\theta},k}(\mathbf{x}) - \bar{\mathbf{s}}_{t,\bar{\boldsymbol{\theta}},k}(\mathbf{x}) \right\|_{L^2(p_t)} = \mathbb{E}_{\mathbf{x} \sim p_t}^{1/2} \left[ \left| \mathbf{s}_{t,\boldsymbol{\theta},k}(\mathbf{x}) - \bar{\mathbf{s}}_{t,\bar{\boldsymbol{\theta}},k}(\mathbf{x}) \right|^2 \right] \\ &= \mathbb{E}_{\mathbf{x} \sim p_t}^{1/2} \left[ \left| \frac{1}{m} \sum_{i=1}^m a_{i,k} \sigma(\mathbf{w}_i^\top \mathbf{x} + \mathbf{u}_i^\top \mathbf{e}(t)) - \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t))] \right|^2 \right]. \end{aligned}$$

If  $(\tilde{\mathbf{W}}, \tilde{\mathbf{U}})$  is different from  $(\mathbf{W}, \mathbf{U})$  at only one component indexed by  $i$ , we have

$$\begin{aligned} & \left| Z_{t,k}(\mathbf{W}, \mathbf{U}) - Z_{t,k}(\tilde{\mathbf{W}}, \tilde{\mathbf{U}}) \right| \\ &= \left| \left\| \mathbf{s}_{t,\boldsymbol{\theta},k}(\mathbf{x}) - \bar{\mathbf{s}}_{t,\bar{\boldsymbol{\theta}},k}(\mathbf{x}) \right\|_{L^2(p_t)} - \left\| \mathbf{s}_{t,\tilde{\boldsymbol{\theta}},k}(\mathbf{x}) - \bar{\mathbf{s}}_{t,\bar{\boldsymbol{\theta}},k}(\mathbf{x}) \right\|_{L^2(p_t)} \right| \\ &\stackrel{(i)}{\leq} \left\| \mathbf{s}_{t,\boldsymbol{\theta},k}(\mathbf{x}) - \mathbf{s}_{t,\tilde{\boldsymbol{\theta}},k}(\mathbf{x}) \right\|_{L^2(p_t)} \\ &= \frac{1}{m} \left\| a_{i,k} \sigma(\mathbf{w}_i^\top \mathbf{x} + \mathbf{u}_i^\top \mathbf{e}(t)) - \tilde{a}_{i,k} \sigma(\tilde{\mathbf{w}}_i^\top \mathbf{x} + \tilde{\mathbf{u}}_i^\top \mathbf{e}(t)) \right\|_{L^2(p_t)} \\ &\leq \frac{1}{m} \left( |a_{i,k}| \left\| \sigma(\mathbf{w}_i^\top \mathbf{x} + \mathbf{u}_i^\top \mathbf{e}(t)) \right\|_{L^2(p_t)} + |\tilde{a}_{i,k}| \left\| \sigma(\tilde{\mathbf{w}}_i^\top \mathbf{x} + \tilde{\mathbf{u}}_i^\top \mathbf{e}(t)) \right\|_{L^2(p_t)} \right) \\ &\stackrel{(ii)}{\leq} \frac{1}{m} (|a_{i,k}| (\|\mathbf{w}_i\|_1 C_{T,\delta} + \|\mathbf{u}_i\|_1 \|\mathbf{e}(t)\|_\infty) + |\tilde{a}_{i,k}| (\|\tilde{\mathbf{w}}_i\|_1 C_{T,\delta} + \|\tilde{\mathbf{u}}_i\|_1 \|\mathbf{e}(t)\|_\infty)) \\ &\stackrel{(iii)}{\leq} \frac{1}{m} (|a_{i,k}| + |\tilde{a}_{i,k}|) (C_{T,\delta} + \|\mathbf{e}(t)\|_\infty) \\ &\stackrel{(iv)}{\lesssim} \frac{1}{m} (C_{T,\delta} + C_{T,e}), \end{aligned}$$

where (i) is from the triangle inequality, (ii) is due to the fact that  $|\sigma(y)| = |\text{ReLU}(y)| \leq |y|$  for any  $y \in \mathbb{R}$ , the triangle and Hölder's inequality and (27), (iii) follows from the positive homogeneity property of the ReLU activation, and (iv) is due to the boundness of the trainable parameter  $\mathbf{a}$  and

embedding function  $e(\cdot)$ . By McDiarmid's inequality (see e.g., Lemma 26.4 in [43]), for any  $\delta > 0$ , with the probability of at least  $1 - \delta$ , we have

$$|Z_{t,k}(\mathbf{W}, \mathbf{U}) - \mathbb{E}_{\mathbf{W}, \mathbf{U}}[Z_{t,k}(\mathbf{W}, \mathbf{U})]| \lesssim \frac{1}{m} (C_{T,\delta} + C_{T,e}) \sqrt{m \log(2/\delta)/2} \quad (29)$$

$$\lesssim (C_{T,\delta} + C_{T,e}) \sqrt{\frac{\log(1/\delta)}{m}}. \quad (30)$$

Since

$$\begin{aligned} & \mathbb{E}_{\mathbf{W}, \mathbf{U}} [Z_{t,k}^2(\mathbf{W}, \mathbf{U})] \\ &= \mathbb{E}_{\mathbf{W}, \mathbf{U}} \left[ \mathbb{E}_{\mathbf{x} \sim p_t} \left[ |s_{t,\theta,k}(\mathbf{x}) - \bar{s}_{t,\bar{\theta},k}(\mathbf{x})|^2 \right] \right] \\ &\stackrel{(v)}{=} \mathbb{E}_{\mathbf{x} \sim p_t} \left[ \mathbb{E}_{\mathbf{W}, \mathbf{U}} \left[ |s_{t,\theta,k}(\mathbf{x}) - \bar{s}_{t,\bar{\theta},k}(\mathbf{x})|^2 \right] \right] \\ &= \frac{1}{m^2} \mathbb{E}_{\mathbf{x} \sim p_t} \left[ \mathbb{E}_{\mathbf{W}, \mathbf{U}} \left[ \left| \sum_{i=1}^m \left( a_{i,k} \sigma(\mathbf{w}_i^\top \mathbf{x} + \mathbf{u}_i^\top \mathbf{e}(t)) - \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t))] \right) \right|^2 \right] \right] \\ &= \frac{1}{m^2} \mathbb{E}_{\mathbf{x} \sim p_t} \left[ \mathbb{E}_{\mathbf{W}, \mathbf{U}} \left[ \sum_{i=1}^m \left( a_{i,k} \sigma(\mathbf{w}_i^\top \mathbf{x} + \mathbf{u}_i^\top \mathbf{e}(t)) - \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t))] \right)^2 \right] \right] \\ &\quad + \frac{1}{m^2} \mathbb{E}_{\mathbf{x} \sim p_t} \left[ \mathbb{E}_{\mathbf{W}, \mathbf{U}} \left[ \sum_{i \neq j} \left( a_{i,k} \sigma(\mathbf{w}_i^\top \mathbf{x} + \mathbf{u}_i^\top \mathbf{e}(t)) - \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t))] \right) \right. \right. \\ &\quad \quad \left. \left. \times \left( a_{j,k} \sigma(\mathbf{w}_j^\top \mathbf{x} + \mathbf{u}_j^\top \mathbf{e}(t)) - \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t))] \right) \right] \right] \\ &= \frac{1}{m^2} \mathbb{E}_{\mathbf{x} \sim p_t} \left[ \sum_{i=1}^m \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} \left[ \left( a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t)) - \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t))] \right)^2 \right] \right] \\ &\quad + \frac{1}{m^2} \mathbb{E}_{\mathbf{x} \sim p_t} \left[ \sum_{i \neq j} \mathbb{E}_{(\mathbf{w}_i, \mathbf{u}_i) \sim \rho_0} \left[ \left( a_{i,k} \sigma(\mathbf{w}_i^\top \mathbf{x} + \mathbf{u}_i^\top \mathbf{e}(t)) - \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t))] \right) \right. \right. \\ &\quad \quad \left. \left. \times \mathbb{E}_{(\mathbf{w}_j, \mathbf{u}_j) \sim \rho_0} \left[ \left( a_{j,k} \sigma(\mathbf{w}_j^\top \mathbf{x} + \mathbf{u}_j^\top \mathbf{e}(t)) - \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t))] \right) \right] \right] \right] \\ &= \frac{1}{m^2} \mathbb{E}_{\mathbf{x} \sim p_t} \left[ \sum_{i=1}^m \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} \left[ \left( a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t)) - \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} [a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t))] \right)^2 \right] \right] \\ &\leq \frac{1}{m} \mathbb{E}_{\mathbf{x} \sim p_t} \left[ \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} \left[ \left( a_k(\mathbf{w}, \mathbf{u}) \sigma(\mathbf{w}^\top \mathbf{x} + \mathbf{u}^\top \mathbf{e}(t)) \right)^2 \right] \right] \\ &\stackrel{(ii)}{\leq} \frac{1}{m} \mathbb{E}_{\mathbf{x} \sim p_t} \left[ \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} \left[ \left( |a_k(\mathbf{w}, \mathbf{u})| (\|\mathbf{w}\|_1 C_{T,\delta} + \|\mathbf{u}\|_1 \|\mathbf{e}(t)\|_\infty) \right)^2 \right] \right] \\ &\stackrel{(iii)}{\leq} \frac{1}{m} \mathbb{E}_{\mathbf{x} \sim p_t} \left[ \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} \left[ \left( |a_k(\mathbf{w}, \mathbf{u})| (C_{T,\delta} + \|\mathbf{e}(t)\|_\infty) \right)^2 \right] \right] \\ &\stackrel{(iv)}{\lesssim} \frac{1}{m} \mathbb{E}_{\mathbf{x} \sim p_t} \left[ \mathbb{E}_{(\mathbf{w}, \mathbf{u}) \sim \rho_0} \left[ (C_{T,\delta} + C_{T,e})^2 \right] \right] \quad (31) \\ &\lesssim (C_{T,\delta} + C_{T,e})^2 \frac{1}{m}, \quad (32) \end{aligned}$$

where (v) is due to Fubini's theorem, and (ii), (iii), (iv) is the same as before. By the triangle inequality, Jensen's inequality, (29) and (31), we obtain

$$\mathbb{E}_{\mathbf{x} \sim p_t} [\|\mathbf{s}_{t,\theta}(\mathbf{x}) - \bar{\mathbf{s}}_{t,\bar{\theta}}(\mathbf{x})\|_2^2] = \sum_{k=1}^d \mathbb{E}_{\mathbf{x} \sim p_t} [|s_{t,\theta,k}(\mathbf{x}) - \bar{s}_{t,\bar{\theta},k}(\mathbf{x})|^2]$$

$$\begin{aligned}
&= \sum_{k=1}^d Z_{t,k}^2(\mathbf{W}, \mathbf{U}) \\
&\leq \sum_{k=1}^d (|Z_{t,k}(\mathbf{W}, \mathbf{U}) - \mathbb{E}_{\mathbf{W}, \mathbf{U}}[Z_{t,k}(\mathbf{W}, \mathbf{U})]| + |\mathbb{E}_{\mathbf{W}, \mathbf{U}}[Z_{t,k}(\mathbf{W}, \mathbf{U})]|)^2 \\
&\lesssim \sum_{k=1}^d (|Z_{t,k}(\mathbf{W}, \mathbf{U}) - \mathbb{E}_{\mathbf{W}, \mathbf{U}}[Z_{t,k}(\mathbf{W}, \mathbf{U})]|^2 + \mathbb{E}_{\mathbf{W}, \mathbf{U}}[Z_{t,k}^2(\mathbf{W}, \mathbf{U})]) \\
&\lesssim \sum_{k=1}^d (C_{T,\delta} + C_{T,e})^2 \frac{\log(1/\delta)}{m} = (C_{T,\delta} + C_{T,e})^2 \frac{\log(1/\delta)}{m} d,
\end{aligned} \tag{33}$$

which gives

$$\begin{aligned}
\text{Err}_{\text{MC}} &= \frac{1}{T} \int_0^T \lambda(t) \cdot \mathbb{E}_{\mathbf{x} \sim p_t} [\|\mathbf{s}_{t,\theta}(\mathbf{x}) - \bar{\mathbf{s}}_{t,\theta}(\mathbf{x})\|_2^2] dt \\
&\lesssim (C_{T,\delta} + C_{T,e})^2 \frac{\log(1/\delta)}{m} d \cdot \frac{1}{T} \int_0^T \lambda(t) dt \leq \frac{\log^2(1/\delta^2)}{m} d.
\end{aligned}$$

Obviously,  $\frac{1}{m} \|\mathbf{A}\|_F^2 = \frac{1}{m} \sum_{i=1}^m \|\mathbf{a}(\mathbf{w}_i, \mathbf{u}_i)\|_2^2 = \|\mathbf{s}_{t,\theta}\|_{\mathcal{H}_{k\rho_0}}^2 < \infty$ . The proof is completed.  $\square$

Now we are ready to prove the main theorem.

*Proof of Theorem 1.* Based on Lemma 2, we have

$$D_{\text{KL}}(p_0 \| p_{0, \hat{\theta}_n(\tau)}) \leq \tilde{\mathcal{L}}(\hat{\theta}_n(\tau); g^2(\cdot)) + D_{\text{KL}}(p_T \| \pi).$$

To bound the term  $\tilde{\mathcal{L}}(\hat{\theta}_n(\tau); g^2(\cdot))$ , we use the following decomposition:

$$\begin{aligned}
\tilde{\mathcal{L}}(\hat{\theta}_n(\tau)) &= \left[ \tilde{\mathcal{L}}(\hat{\theta}_n(\tau)) - \tilde{\mathcal{L}}(\theta(\tau)) \right] + \tilde{\mathcal{L}}(\theta(\tau)) \\
&\lesssim \left[ \tilde{\mathcal{L}}(\hat{\theta}_n(\tau)) - \tilde{\mathcal{L}}(\theta(\tau)) \right] + \tilde{\mathcal{L}}(\bar{\theta}(\tau)) + \text{Monte Carlo} \triangleq I_1 + I_2 + I_3.
\end{aligned}$$

According to Lemma 4, we obtain

$$I_2 \lesssim \tilde{\mathcal{L}}(\bar{\theta}^*) + \frac{1}{\tau} \left( \|\bar{\mathbf{s}}_{0, \bar{\theta}_0}\|_{\mathcal{H}}^2 + \|\bar{\mathbf{s}}_{0, \bar{\theta}^*}\|_{\mathcal{H}}^2 \right).$$

The term  $I_3$  can be further divided into

$$\begin{aligned}
I_3 &:= \mathbb{E}_{t \sim \mathcal{U}(0, T)} \left[ \lambda(t) \cdot \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \|\mathbf{s}_{t, \theta(\tau)}(\mathbf{x}(t)) - \bar{\mathbf{s}}_{t, \bar{\theta}(\tau)}(\mathbf{x}(t))\|_2^2 \right] \right] \\
&\lesssim \mathbb{E}_{t \sim \mathcal{U}(0, T)} \left[ \lambda(t) \cdot \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \|\bar{\mathbf{s}}_{t, \theta(\tau)}(\mathbf{x}(t)) - \bar{\mathbf{s}}_{t, \bar{\theta}^*}(\mathbf{x}(t))\|_2^2 \right] \right] \\
&\quad + \mathbb{E}_{t \sim \mathcal{U}(0, T)} \left[ \lambda(t) \cdot \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \|\mathbf{s}_{t, \theta^*}(\mathbf{x}(t)) - \bar{\mathbf{s}}_{t, \bar{\theta}^*}(\mathbf{x}(t))\|_2^2 \right] \right] \\
&\quad + \mathbb{E}_{t \sim \mathcal{U}(0, T)} \left[ \lambda(t) \cdot \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \|\mathbf{s}_{t, \theta(\tau)}(\mathbf{x}(t)) - \mathbf{s}_{t, \theta^*}(\mathbf{x}(t))\|_2^2 \right] \right] \\
&=: I_{3,1} + I_{3,2} + I_{3,3},
\end{aligned}$$

where  $\theta^*$  is the Monte Carlo estimator of  $\bar{\theta}^*$ . By Lemma 6, for any  $\delta > 0$ ,  $\delta \ll 1$ , with the probability of at least  $1 - \delta$ , it holds that

$$I_{3,2} = \text{Err}_{\text{MC}}(\theta^*, \bar{\theta}^*; T, \lambda(\cdot)) \lesssim \frac{\log^2(1/\delta^2)}{m} d, \tag{34}$$

while Lemma 4 gives

$$I_{3,1} \lesssim \tilde{\mathcal{L}}(\bar{\theta}(\tau)) + \tilde{\mathcal{L}}(\bar{\theta}^*) = I_2 + \tilde{\mathcal{L}}(\bar{\theta}^*) \lesssim \tilde{\mathcal{L}}(\bar{\theta}^*) + \frac{1}{\tau} \left( \|\bar{\mathbf{s}}_{0, \bar{\theta}_0}\|_{\mathcal{H}}^2 + \|\bar{\mathbf{s}}_{0, \bar{\theta}^*}\|_{\mathcal{H}}^2 \right),$$

and similarly,

$$I_{3,3} \lesssim \tilde{\mathcal{L}}(\boldsymbol{\theta}^*) + \frac{1}{\tau} \left( \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}}^2 + \|\mathbf{s}_{0,\boldsymbol{\theta}^*}\|_{\mathcal{H}}^2 \right).$$

Hence, for any  $\delta > 0$ ,  $\delta \ll 1$ , with the probability of at least  $1 - \delta$ , it holds that

$$I_3 \lesssim \frac{\log^2(1/\delta^2)}{m} d + \tilde{\mathcal{L}}(\bar{\boldsymbol{\theta}}^*) + \tilde{\mathcal{L}}(\boldsymbol{\theta}^*) + \frac{1}{\tau} \left( \|\bar{\mathbf{s}}_{0,\bar{\boldsymbol{\theta}}_0}\|_{\mathcal{H}}^2 + \|\bar{\mathbf{s}}_{0,\bar{\boldsymbol{\theta}}^*}\|_{\mathcal{H}}^2 + \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}}^2 + \|\mathbf{s}_{0,\boldsymbol{\theta}^*}\|_{\mathcal{H}}^2 \right).$$

For the term  $I_1$ , we have

$$\begin{aligned} & \sqrt{\tilde{\mathcal{L}}(\hat{\boldsymbol{\theta}}_n(\tau))} - \sqrt{\tilde{\mathcal{L}}(\boldsymbol{\theta}(\tau))} \\ &= \left\{ \mathbb{E}_{t \sim \mathcal{U}(0,T)} \left[ \lambda(t) \cdot \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \|\mathbf{s}_{t,\hat{\boldsymbol{\theta}}_n(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t))\|_2^2 \right] \right] \right\}^{\frac{1}{2}} \\ & \quad - \left\{ \mathbb{E}_{t \sim \mathcal{U}(0,T)} \left[ \lambda(t) \cdot \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \|\mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t))\|_2^2 \right] \right] \right\}^{\frac{1}{2}} \\ &= \left\{ \mathbb{E}_{t \sim \mathcal{U}(0,T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \|\mathbf{s}_{t,\hat{\boldsymbol{\theta}}_n(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t))\|_2^2 \right] \right\}^{\frac{1}{2}} \\ & \quad - \left\{ \mathbb{E}_{t \sim \mathcal{U}(0,T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \|\mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t))\|_2^2 \right] \right\}^{\frac{1}{2}} \\ &\leq \left\{ \mathbb{E}_{t \sim \mathcal{U}(0,T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \|\mathbf{s}_{t,\hat{\boldsymbol{\theta}}_n(\tau)}(\mathbf{x}(t)) - \mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t))\|_2^2 \right] \right\}^{\frac{1}{2}} \\ &= \left\{ \mathbb{E}_{t \sim \mathcal{U}(0,T)} \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \lambda(t) \left\| \frac{1}{m} \sum_{i=1}^m \hat{\mathbf{a}}_i(\tau) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) - \frac{1}{m} \sum_{i=1}^m \mathbf{a}_i(\tau) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right\|_2^2 \right] \right\}^{\frac{1}{2}}. \end{aligned}$$

Notice that

$$\begin{aligned} & \left\| \frac{1}{m} \sum_{i=1}^m (\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right\|_2^2 \\ &\leq \frac{1}{m^2} \left( \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2 |\sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t))| \right)^2 \\ &\leq \frac{1}{m^2} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 \sum_{i=1}^m |\sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t))|^2 \\ &\lesssim \frac{1}{m^2} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 \sum_{i=1}^m \left( |\mathbf{w}_i^\top \mathbf{x}(t)|^2 + |\mathbf{u}_i^\top \mathbf{e}(t)|^2 \right) \\ &\leq \frac{1}{m^2} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 \sum_{i=1}^m \left( \|\mathbf{w}_i\|_1^2 \|\mathbf{x}(t)\|_\infty^2 + \|\mathbf{u}_i\|_1^2 \|\mathbf{e}(t)\|_\infty^2 \right) \\ &\lesssim \frac{1}{m^2} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 \sum_{i=1}^m (C_{T,\delta}^2 + C_{T,e}^2), \end{aligned}$$

which gives

$$\begin{aligned} \sqrt{\tilde{\mathcal{L}}(\hat{\boldsymbol{\theta}}_n(\tau))} - \sqrt{\tilde{\mathcal{L}}(\boldsymbol{\theta}(\tau))} &\lesssim \left\{ \frac{1}{m} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 (C_{T,\delta}^2 + C_{T,e}^2) \right\}^{\frac{1}{2}} \\ &\leq (C_{T,\delta} + C_{T,e}) \left\{ \frac{1}{m} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 \right\}^{\frac{1}{2}}. \end{aligned}$$

Here, the triangle inequality, Cauchy–Schwartz inequality, the fact that  $|\sigma(y)| = |\text{ReLU}(y)| \leq |y|$  for any  $y \in \mathbb{R}$ , Hölder’s inequality, the positive homogeneity property of the ReLU activation, and

the boundness of the input data, embedding function  $e(t)$  and weighting function  $\lambda(t)$ . Thus, we get

$$\begin{aligned} & \tilde{\mathcal{L}}(\hat{\boldsymbol{\theta}}_n(\tau)) - \tilde{\mathcal{L}}(\boldsymbol{\theta}(\tau)) \\ & \lesssim \frac{1}{m} (C_{T,\delta}^2 + C_{T,e}^2) \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 + \sqrt{\tilde{\mathcal{L}}(\boldsymbol{\theta}(\tau))} (C_{T,\delta} + C_{T,e}) \left\{ \frac{1}{m} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 \right\}^{\frac{1}{2}} \\ & \lesssim \sqrt{\tilde{\mathcal{L}}(\boldsymbol{\theta}^*)} + \frac{1}{\tau} \left( \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}}^2 + \|\mathbf{s}_{0,\boldsymbol{\theta}^*}\|_{\mathcal{H}}^2 \right) (C_{T,\delta} + C_{T,e}) \left\{ \frac{1}{m} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 \right\}^{\frac{1}{2}} \\ & \quad + \frac{1}{m} (C_{T,\delta}^2 + C_{T,e}^2) \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2, \end{aligned}$$

where the last inequality follows from Lemma 4. We further deduce that

$$\begin{aligned} & \frac{1}{m} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 \\ & = \frac{1}{m} \sum_{i=1}^m \left\| \int_0^\tau \frac{d}{d\tau_0} (\hat{\mathbf{a}}_i(\tau_0) - \mathbf{a}_i(\tau_0)) d\tau_0 \right\|_2^2 \\ & = \frac{1}{m} \sum_{i=1}^m \left\| \int_0^\tau \left( \nabla_{\boldsymbol{\theta}_i(\tau_0)} \tilde{\mathcal{L}}(\boldsymbol{\theta}(\tau_0)) - \nabla_{\hat{\boldsymbol{\theta}}_{n,i}(\tau_0)} \hat{\mathcal{L}}_n(\hat{\boldsymbol{\theta}}_n(\tau_0)) \right) d\tau_0 \right\|_2^2 \\ & = \frac{1}{m^2} \sum_{i=1}^m \left\| \int_0^\tau \left( 2\mathbb{E}_{t \sim \mathcal{U}(0,T)} \left[ \lambda(t) \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \left( \mathbf{s}_{t,\boldsymbol{\theta}(\tau_0)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t)) \right) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right] \right] \right. \right. \\ & \quad \left. \left. - 2\mathbb{E}_{t \sim \mathcal{U}(0,T)} \left[ \lambda(t) \mathbb{E}_{\mathbf{x}(t) \sim p_t^{(n)}} \left[ \left( \mathbf{s}_{t,\hat{\boldsymbol{\theta}}_n(\tau_0)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t)) \right) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right] \right] \right] \right) d\tau_0 \right\|_2^2, \end{aligned}$$

where  $p_t^{(n)}$  denotes the empirical distribution of  $p_t$ . Note that

$$\|\boldsymbol{\theta}\|_2^2 = \|\text{vec}(\mathbf{A})\|_2^2/m = \|\mathbf{A}\|_F^2/m = \|\mathbf{s}_{0,\boldsymbol{\theta}}\|_{\mathcal{H}}^2, \quad (35)$$

by Lemma 5 we get  $\|\boldsymbol{\theta}(\tau)\|_2 = \|\mathbf{s}_{0,\boldsymbol{\theta}(\tau)}\|_{\mathcal{H}} \lesssim \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m}$ . For any  $t \in [0, T]$ , define the function space

$$\mathcal{F}_t := \{ \mathbf{f}_1(\mathbf{x}(t); \boldsymbol{\theta}_1(\tau)) \mathbf{f}_2(\mathbf{x}(t); \boldsymbol{\theta}_2) : \mathbf{f}_1 \in \mathcal{F}_{1,t}, \mathbf{f}_2 \in \mathcal{F}_{2,t} \},$$

where

$$\begin{aligned} \mathcal{F}_{1,t} & := \left\{ \mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t)) : \|\boldsymbol{\theta}(\tau)\|_2 \lesssim \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m} \right\}, \\ \mathcal{F}_{2,t} & := \left\{ \sigma(\mathbf{w}^\top \mathbf{x}(t) + \mathbf{u}^\top \mathbf{e}(t)) : \|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1 \right\}. \end{aligned}$$

Then, according to Theorem A.5 in [61], for any  $\delta \in (0, 1)$ , with the probability at least  $1 - \delta$  over the choice of the dataset  $\mathcal{D}_x = \{\mathbf{x}_i\}_{i=1}^n$ , it holds that

$$\begin{aligned} & \left| \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \left( \mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t)) \right) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right] \right. \\ & \quad \left. - \mathbb{E}_{\mathbf{x}(t) \sim p_t^{(n)}} \left[ \left( \mathbf{s}_{t,\hat{\boldsymbol{\theta}}_n(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t)) \right) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right] \right| \\ & \leq \left| \mathbb{E}_{\mathbf{x}(t) \sim p_t} \left[ \left( \mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t)) \right) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right] \right. \\ & \quad \left. - \mathbb{E}_{\mathbf{x}(t) \sim p_t^{(n)}} \left[ \left( \mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t)) \right) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right] \right| \\ & \quad + \left| \mathbb{E}_{\mathbf{x}(t) \sim p_t^{(n)}} \left[ \left( \mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t)) \right) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right] \right. \\ & \quad \left. - \mathbb{E}_{\mathbf{x}(t) \sim p_t^{(n)}} \left[ \left( \mathbf{s}_{t,\hat{\boldsymbol{\theta}}_n(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t)) \right) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right] \right| \\ & \lesssim \widehat{\text{Rad}}_n(\mathcal{F}_t) + \sup_{\mathbf{f} \in \mathcal{F}_t, \mathbf{x}(t) \in [-C_{T,\delta}, C_{T,\delta}]^d} |\mathbf{f}(\mathbf{x}(t))| \sqrt{\frac{\log(2/\delta)}{n}} \\ & \quad + \left| \mathbb{E}_{\mathbf{x}(t) \sim p_t^{(n)}} \left[ \left( \mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) - \mathbf{s}_{t,\hat{\boldsymbol{\theta}}_n(\tau)}(\mathbf{x}(t)) \right) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right] \right| =: J_1 + J_2 + J_3, \end{aligned}$$

where  $\widehat{\text{Rad}}_n$  denotes the (empirical) Rademacher complexity of  $\mathcal{F}_t$  on  $D_{\mathbf{x}} = \{\mathbf{x}_i\}_{i=1}^n$ , and all the inequalities hold in the element-wise sense.

(i) For  $J_1$ , according to Lemma A.6 in [61], we have

$$\widehat{\text{Rad}}_n(\mathcal{F}_t) \leq \left( \sup_{\mathbf{f}_1 \in \mathcal{F}_{1,t}, \mathbf{x}(t) \in [-C_{T,\delta}, C_{T,\delta}]^d} |\mathbf{f}_1(\mathbf{x}(t))| + \sup_{\mathbf{f}_2 \in \mathcal{F}_{2,t}, \mathbf{x}(t) \in [-C_{T,\delta}, C_{T,\delta}]^d} |f_2(\mathbf{x}(t))| \right) \cdot \left( \widehat{\text{Rad}}_n(\mathcal{F}_{1,t}) + \widehat{\text{Rad}}_n(\mathcal{F}_{2,t}) \right).$$

Note that

$$\begin{aligned} |\sigma(\mathbf{w}^\top \mathbf{x}(t) + \mathbf{u}^\top \mathbf{e}(t))| &\leq |\mathbf{w}^\top \mathbf{x}(t) + \mathbf{u}^\top \mathbf{e}(t)| \\ &\lesssim \|\mathbf{w}\|_1 \|\mathbf{x}(t)\|_\infty + \|\mathbf{u}\|_1 \|\mathbf{e}(t)\|_\infty \\ &\leq C_{T,\delta} + C_{T,e}, \end{aligned} \quad (36)$$

which yields

$$\begin{aligned} |\mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t))| &= \left| \frac{1}{\sqrt{m}} \sum_{i=1}^m \boldsymbol{\theta}_i(\tau) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right| \\ &\leq \frac{1}{\sqrt{m}} \sum_{i=1}^m |\boldsymbol{\theta}_i(\tau)| |\sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t))| \\ &\leq \frac{1}{\sqrt{m}} (C_{T,\delta} + C_{T,e}) \sum_{i=1}^m |\boldsymbol{\theta}_i(\tau)| \\ &\leq (C_{T,\delta} + C_{T,e}) \|\boldsymbol{\theta}(\tau)\|_2 \\ &\lesssim (C_{T,\delta} + C_{T,e}) \left( \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m} \right), \end{aligned} \quad (37)$$

and hence

$$\begin{aligned} |\mathbf{f}_1(\mathbf{x}(t))| &= |\mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t))| \\ &\lesssim (C_{T,\delta} + C_{T,e}) \left( \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m} \right) + C'_{T,\delta}, \\ |f_2(\mathbf{x}(t))| &= |\sigma(\mathbf{w}^\top \mathbf{x}(t) + \mathbf{u}^\top \mathbf{e}(t))| \lesssim C_{T,\delta} + C_{T,e}, \end{aligned}$$

where  $C'_{T,\delta} := \max_{\mathbf{x}(t) \in [-C_{T,\delta}, C_{T,\delta}]^d} |\nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t))|$ . This gives

$$\widehat{\text{Rad}}_n(\mathcal{F}_t) \lesssim (C_{T,\delta} + C_{T,e}) \left( \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m} + 1 \right) \left( \widehat{\text{Rad}}_n(\mathcal{F}_{1,t}) + \widehat{\text{Rad}}_n(\mathcal{F}_{2,t}) \right).$$

Let

$$\mathcal{F}'_{1,t} := \left\{ \mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) : \|\boldsymbol{\theta}(\tau)\|_2 \lesssim \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m} \right\},$$

according to Lemma 26.6 in [43], we get  $\widehat{\text{Rad}}_n(\mathcal{F}_{1,t}) \leq \widehat{\text{Rad}}_n(\mathcal{F}'_{1,t})$ . Since

$$\begin{aligned} n\widehat{\text{Rad}}_n(\mathcal{F}'_{1,t}) &= \mathbb{E}_{\boldsymbol{\xi}} \left[ \sup_{\|\boldsymbol{\theta}(\tau)\|_2 \lesssim \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m}} \sum_{j=1}^n \xi_j \mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}_j(t)) \right] \\ &= \mathbb{E}_{\boldsymbol{\xi}} \left[ \sup_{\|\boldsymbol{\theta}(\tau)\|_2 \lesssim \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m}} \sum_{j=1}^n \xi_j \frac{1}{\sqrt{m}} \sum_{i=1}^m \boldsymbol{\theta}_i(\tau) \sigma(\mathbf{w}_i^\top \mathbf{x}_j(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right] \\ &\leq \frac{1}{\sqrt{m}} \mathbb{E}_{\boldsymbol{\xi}} \left[ \sup_{\substack{\|\boldsymbol{\theta}(\tau)\|_2 \lesssim \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m} \\ \|\mathbf{w}_i\|_1 + \|\mathbf{u}_i\|_1 \leq 1, i \in [m]}} \sum_{i=1}^m \boldsymbol{\theta}_i(\tau) \sum_{j=1}^n \xi_j \sigma(\mathbf{w}_i^\top \mathbf{x}_j(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right] \end{aligned}$$

$$\begin{aligned}
&\leq \frac{1}{\sqrt{m}} \mathbb{E}_{\xi} \left[ \sup_{\|\boldsymbol{\theta}(\tau)\|_2 \lesssim \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m}} \sum_{i=1}^m |\boldsymbol{\theta}_i(\tau)| \sup_{\|\mathbf{w}_i\|_1 + \|\mathbf{u}_i\|_1 \leq 1, i \in [n]} \left| \sum_{j=1}^n \xi_j \sigma(\mathbf{w}_i^\top \mathbf{x}_j(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right| \right] \\
&\leq \mathbb{E}_{\xi} \left[ \sup_{\|\boldsymbol{\theta}(\tau)\|_2 \lesssim \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m}} \|\boldsymbol{\theta}(\tau)\|_2 \sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \left| \sum_{j=1}^n \xi_j \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)) \right| \right] \\
&\lesssim \left( \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m} \right) \mathbb{E}_{\xi} \left[ \sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \left| \sum_{j=1}^n \xi_j \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)) \right| \right],
\end{aligned}$$

where the  $\{\xi_i\}_{i=1}^n$  are independent random variables with the distribution  $\mathbb{P}(\xi_i = 1) = \mathbb{P}(\xi_i = -1) = 1/2$ , and obviously,

$$\sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \sum_{j=1}^n \xi_j \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)) \geq \sum_{j=1}^n \xi_j \sigma(\mathbf{0}_d^\top \mathbf{x}_j(t) + \mathbf{0}_d^\top \mathbf{e}(t)) = 0,$$

we have

$$\begin{aligned}
&\sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \left| \sum_{j=1}^n \xi_j \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)) \right| \\
&\leq \max \left\{ \sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \sum_{j=1}^n \xi_j \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)), \sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \sum_{j=1}^n (-\xi_j) \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)) \right\} \\
&\leq \sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \sum_{j=1}^n \xi_j \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)) + \sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \sum_{j=1}^n (-\xi_j) \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)),
\end{aligned}$$

and by symmetry

$$\begin{aligned}
&\mathbb{E}_{\xi} \left[ \sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \left| \sum_{j=1}^n \xi_j \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)) \right| \right] \\
&\leq \mathbb{E}_{\xi} \left[ \sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \sum_{j=1}^n \xi_j \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)) \right] + \mathbb{E}_{\xi} \left[ \sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \sum_{j=1}^n (-\xi_j) \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)) \right] \\
&= 2 \mathbb{E}_{\xi} \left[ \sup_{\|\mathbf{w}\|_1 + \|\mathbf{u}\|_1 \leq 1} \sum_{j=1}^n \xi_j \sigma(\mathbf{w}^\top \mathbf{x}_j(t) + \mathbf{u}^\top \mathbf{e}(t)) \right] = 2n \widehat{\text{Rad}}_n(\mathcal{F}_{2,t}),
\end{aligned}$$

i.e.,  $\widehat{\text{Rad}}_n(\mathcal{F}_{1,t}) \leq \widehat{\text{Rad}}_n(\mathcal{F}'_{1,t}) \lesssim 2 \left( \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m} \right) \widehat{\text{Rad}}_n(\mathcal{F}_{2,t})$ . According to Lemma 26.9 (contraction lemma) and Lemma 26.11 in [43], we have

$$\widehat{\text{Rad}}_n(\mathcal{F}_{2,t}) \leq (\|\mathbf{x}(t)\|_{\infty} + \|\mathbf{e}(t)\|_{\infty}) \sqrt{\frac{2 \log(4d)}{n}} \lesssim (C_{T,\delta} + C_{T,e}) \sqrt{\frac{\log(d+1)}{n}}.$$

Combining above, we obtain

$$J_1 = \widehat{\text{Rad}}_n(\mathcal{F}_t) \lesssim (C_{T,\delta} + C_{T,e})^2 \left( \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m} + 1 \right)^2 \sqrt{\frac{\log(d+1)}{n}}.$$

(ii) For  $J_2$ , by (36) and (37), we get

$$\begin{aligned}
|\mathbf{f}(\mathbf{x}(t))| &= |\mathbf{s}_{t,\boldsymbol{\theta}(\tau)}(\mathbf{x}(t)) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t))| |\sigma(\mathbf{w}^\top \mathbf{x}(t) + \mathbf{u}^\top \mathbf{e}(t))| \\
&\lesssim (C_{T,\delta} + C_{T,e})^2 \left( \|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}} + \sqrt{\tau/m} \right) + C'_{T,\delta} (C_{T,\delta} + C_{T,e}),
\end{aligned}$$

which gives

$$J_2 \lesssim (C_{T,\delta} + C_{T,e})^2 \left( \|\mathbf{s}_{0,\theta_0}\|_{\mathcal{H}} + \sqrt{\tau/m} + 1 \right) \sqrt{\frac{\log(1/\delta)}{n}}.$$

(iii) For  $J_3$ , we similarly have

$$\left| \mathbf{s}_{t,\hat{\theta}_n(\tau)}(\mathbf{x}(t)) \right| \lesssim (C_{T,\delta} + C_{T,e}) \left( \|\mathbf{s}_{0,\theta_0}\|_{\mathcal{H}} + \sqrt{\tau/m} \right),$$

hence by (36) and (37), we get

$$\begin{aligned} J_3 &\leq \mathbb{E}_{\mathbf{x}(t) \sim p_t^{(n)}} \left| \left( \mathbf{s}_{t,\theta(\tau)}(\mathbf{x}(t)) - \mathbf{s}_{t,\hat{\theta}_n(\tau)}(\mathbf{x}(t)) \right) \sigma(\mathbf{w}_i^\top \mathbf{x}(t) + \mathbf{u}_i^\top \mathbf{e}(t)) \right| \\ &\lesssim (C_{T,\delta} + C_{T,e})^2 \left( \|\mathbf{s}_{0,\theta_0}\|_{\mathcal{H}} + \sqrt{\tau/m} \right). \end{aligned}$$

Combining (i), (ii) and (iii), we obtain

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 &\lesssim \frac{1}{m^2} \sum_{i=1}^m \left\| \int_0^\tau \mathbb{E}_{t \sim \mathcal{U}(0,T)} [\lambda(t)(J_1 + J_2 + J_3)\mathbf{1}_d] d\tau_0 \right\|_2^2 \\ &\lesssim \frac{1}{m^2} \sum_{i=1}^m \left\| \int_0^\tau (J_1 + J_2 + J_3)\mathbf{1}_d d\tau_0 \right\|_2^2 \\ &= (J_1 + J_2 + J_3)^2 \tau^2 \frac{d}{m} \\ &\lesssim \tau^2 \frac{d}{m} (C_{T,\delta} + C_{T,e})^4 \left[ \left( \|\mathbf{s}_{0,\theta_0}\|_{\mathcal{H}}^2 + \frac{\tau}{m} + 1 \right)^2 \right. \\ &\quad \left. \cdot \left( \sqrt{\frac{\log(d+1)}{n}} + \sqrt{\frac{\log(1/\delta)}{n}} \right)^2 + \left( \|\mathbf{s}_{0,\theta_0}\|_{\mathcal{H}}^2 + \frac{\tau}{m} \right) \right], \end{aligned}$$

which gives

$$\begin{aligned} I_1 &= \tilde{\mathcal{L}}(\hat{\theta}_n(\tau)) - \tilde{\mathcal{L}}(\theta(\tau)) \\ &\lesssim (C_{T,\delta} + C_{T,e}) \left( \sqrt{\tilde{\mathcal{L}}(\theta^*)} + \frac{1}{\sqrt{\tau}} (\|\mathbf{s}_{0,\theta_0}\|_{\mathcal{H}} + \|\mathbf{s}_{0,\theta^*}\|_{\mathcal{H}}) \right) \left\{ \frac{1}{m} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 \right\}^{\frac{1}{2}} \\ &\quad + (C_{T,\delta}^2 + C_{T,e}^2) \frac{1}{m} \sum_{i=1}^m \|\hat{\mathbf{a}}_i(\tau) - \mathbf{a}_i(\tau)\|_2^2 \\ &\lesssim (C_{T,\delta} + C_{T,e})^3 \left( \sqrt{\tilde{\mathcal{L}}(\theta^*)} + \frac{1}{\sqrt{\tau}} (\|\mathbf{s}_{0,\theta_0}\|_{\mathcal{H}} + \|\mathbf{s}_{0,\theta^*}\|_{\mathcal{H}}) \right) \\ &\quad \cdot \tau \sqrt{\frac{d}{m}} \left[ \left( \|\mathbf{s}_{0,\theta_0}\|_{\mathcal{H}}^2 + \tau + 1 \right) \left( \sqrt{\frac{\log(d+1)}{n}} + \sqrt{\frac{\log(1/\delta)}{n}} \right) + \left( \|\mathbf{s}_{0,\theta_0}\|_{\mathcal{H}} + \sqrt{\frac{\tau}{m}} \right) \right] \\ &\quad + \tau^2 \frac{d}{m} (C_{T,\delta} + C_{T,e})^6 \left[ \left( \|\mathbf{s}_{0,\theta_0}\|_{\mathcal{H}}^2 + \tau + 1 \right)^2 \left( \sqrt{\frac{\log(d+1)}{n}} + \sqrt{\frac{\log(1/\delta)}{n}} \right)^2 + \left( \|\mathbf{s}_{0,\theta_0}\|_{\mathcal{H}}^2 + \frac{\tau}{m} \right) \right] \\ &\lesssim (C_{T,\delta} + C_{T,e})^6 \tau \sqrt{\frac{d}{m}} \left[ (\tau + 1) \left( \sqrt{\frac{\log(d+1)}{n}} + \sqrt{\frac{\log(1/\delta)}{n}} \right) + \left( \frac{\|\mathbf{A}_0\|_F}{\sqrt{m}} + \sqrt{\frac{\tau}{m}} \right) \right] \\ &\quad \cdot \left\{ \left( \sqrt{\tilde{\mathcal{L}}(\theta^*)} + \frac{1}{\sqrt{m\tau}} (\|\mathbf{A}_0\|_F + \|\mathbf{A}_*\|_F) \right) \right. \\ &\quad \left. + \tau \sqrt{\frac{d}{m}} \left[ (\tau + 1) \left( \sqrt{\frac{\log(d+1)}{n}} + \sqrt{\frac{\log(1/\delta)}{n}} \right) + \left( \frac{\|\mathbf{A}_0\|_F}{\sqrt{m}} + \sqrt{\frac{\tau}{m}} \right) \right] \right\} \quad (38) \end{aligned}$$

$$\lesssim \log^3(1/\delta^2)\tau\sqrt{\frac{d}{m}}\left[(\tau+1)\left(\sqrt{\frac{\log(d+1)}{n}}+\sqrt{\frac{\log(1/\delta)}{n}}\right)+\left(\frac{1}{\sqrt{m}}+\sqrt{\frac{\tau}{m}}\right)\right] \\ \cdot \left\{\left(\sqrt{\tilde{\mathcal{L}}(\boldsymbol{\theta}^*)}+\frac{1}{\sqrt{m\tau}}\right)+\tau\sqrt{\frac{d}{m}}\left[(\tau+1)\left(\sqrt{\frac{\log(d+1)}{n}}+\sqrt{\frac{\log(1/\delta)}{n}}\right)+\left(\frac{1}{\sqrt{m}}+\sqrt{\frac{\tau}{m}}\right)\right]\right\},$$

where  $\lesssim$  hides universal positive constants only depending on  $T$ .

Combining all above estimates yields

$$D_{\text{KL}}\left(p_0\|p_{0,\hat{\boldsymbol{\theta}}_n(\tau)}\right) \\ \lesssim I_1+I_2+I_3+D_{\text{KL}}(p_T\|\pi) \\ \lesssim \log^3(1/\delta^2)\tau\sqrt{\frac{d}{m}}\left[(\tau+1)\left(\sqrt{\frac{\log(d+1)}{n}}+\sqrt{\frac{\log(1/\delta)}{n}}\right)+\frac{1}{\sqrt{m}}(1+\sqrt{\tau})\right] \\ \cdot \left\{\left(\sqrt{\tilde{\mathcal{L}}(\boldsymbol{\theta}^*)}+\frac{1}{\sqrt{m\tau}}\right)+\tau\sqrt{\frac{d}{m}}\left[(\tau+1)\left(\sqrt{\frac{\log(d+1)}{n}}+\sqrt{\frac{\log(1/\delta)}{n}}\right)+\frac{1}{\sqrt{m}}(1+\sqrt{\tau})\right]\right\} \\ +\left(\tilde{\mathcal{L}}(\bar{\boldsymbol{\theta}}^*)+\tilde{\mathcal{L}}(\boldsymbol{\theta}^*)\right)+\frac{\log^2(1/\delta^2)}{m}d+\frac{1}{\tau}\left(\|\bar{\mathbf{s}}_{0,\bar{\boldsymbol{\theta}}_0}\|_{\mathcal{H}}^2+\|\bar{\mathbf{s}}_{0,\bar{\boldsymbol{\theta}}^*}\|_{\mathcal{H}}^2+\|\mathbf{s}_{0,\boldsymbol{\theta}_0}\|_{\mathcal{H}}^2+\|\mathbf{s}_{0,\boldsymbol{\theta}^*}\|_{\mathcal{H}}^2\right)+D_{\text{KL}}(p_T\|\pi).$$

If one only focuses on the  $(m, n)$ -dependence, i.e. the dependence on the model capacity and sample size, this upper bound can be further simplified as

$$D_{\text{KL}}\left(p_0\|p_{0,\hat{\boldsymbol{\theta}}_n(\tau)}\right) \\ \stackrel{\text{(vi)}}{\lesssim} \left(\frac{\tau^2}{\sqrt{mn}}+\frac{\tau\sqrt{\tau}}{m}\right)\cdot\left(\sqrt{\tilde{\mathcal{L}}(\boldsymbol{\theta}^*)}+\frac{1}{\sqrt{m\tau}}+\frac{\tau^2}{\sqrt{mn}}+\frac{\tau\sqrt{\tau}}{m}\right) \\ +\left(\tilde{\mathcal{L}}(\bar{\boldsymbol{\theta}}^*)+\tilde{\mathcal{L}}(\boldsymbol{\theta}^*)\right)+\frac{1}{m}+\frac{1}{\tau}+D_{\text{KL}}(p_T\|\pi) \\ \leq\left(\sqrt{\tilde{\mathcal{L}}(\boldsymbol{\theta}^*)}+\frac{1}{\sqrt{m\tau}}+\frac{\tau^2}{\sqrt{mn}}+\frac{\tau\sqrt{\tau}}{m}\right)^2+\left(\tilde{\mathcal{L}}(\bar{\boldsymbol{\theta}}^*)+\tilde{\mathcal{L}}(\boldsymbol{\theta}^*)\right)+\frac{1}{m}+\frac{1}{\tau}+D_{\text{KL}}(p_T\|\pi) \\ \lesssim\left(\tilde{\mathcal{L}}(\boldsymbol{\theta}^*)+\frac{1}{m\tau}+\frac{\tau^4}{mn}+\frac{\tau^3}{m^2}\right)+\left(\tilde{\mathcal{L}}(\bar{\boldsymbol{\theta}}^*)+\tilde{\mathcal{L}}(\boldsymbol{\theta}^*)\right)+\frac{1}{m}+\frac{1}{\tau}+D_{\text{KL}}(p_T\|\pi) \\ \lesssim\frac{\tau^4}{mn}+\frac{\tau^3}{m^2}+\frac{1}{\tau}+\frac{1}{m}+\left(\tilde{\mathcal{L}}(\bar{\boldsymbol{\theta}}^*)+\tilde{\mathcal{L}}(\boldsymbol{\theta}^*)\right)+D_{\text{KL}}(p_T\|\pi),$$

where we assume  $\tau \geq 1$  for simplicity, and  $\stackrel{\text{(vi)}}{\lesssim}$  hides the term  $d \log(d+1)$ , the polynomials of  $\log(1/\delta^2)$  and finite RKHS norms  $\|\cdot\|_{\mathcal{H}}$ . The proof is completed.  $\square$

**Approximation errors.** For the (universal) approximation, we discuss in two points:

- First, the approximation is a separate problem that can be analyzed independently of training and generalization. While it is beyond the scope of current work, we have included the approximation error in the final estimates. When the approximation by random feature models fails, the generalization error is supposed to be significant.
- In addition, the random feature model can approximate Lipschitz continuous functions on a compact domain (Theorem 6 in [17]). Notice that the forward diffusion process defines a random path  $(\mathbf{x}(t), t)_{t \in [0, T]}$  contained in a rectangular domain  $R_{T, \delta} := [-C_{T, \delta}, C_{T, \delta}]^d \times [0, T] \subset \mathbb{R}^{d+1}$  with  $C_{T, \delta} := C_T(C_{\mathbf{x}} + \sqrt{\log(1/\delta^2)})$  (use Lemma 1 and the boundness of inputs), one can apply Theorem 6 in [17] to bound (7) on the domain  $R_{T, \delta}$  in  $\mathbb{R}^{d+1}$  to obtain approximation results for Lipschitz continuous target score functions.

## A.2 Data-Dependent Generalization Gap

**Lemma 7** (Forward perturbation estimates, Gaussian mixtures). *Suppose that  $x$  is sampled from a one-dimensional 2-mode Gaussian mixture:  $p_0(x) = q_1\mathcal{N}(x; -\mu, 1) + q_2\mathcal{N}(x; \mu, 1)$ , where  $\mu > 0$ ,  $q_1, q_2 > 0$  with  $q_1 + q_2 = 1$  are all constants. Then for any  $\delta > 0$ ,  $\delta \ll 1$ ,  $\mu > \sqrt{\log(1/\delta^2)}$ , with the probability of at least  $1 - \delta$ , we have*

$$|x(t)| \lesssim C_T \left( \mu + \sqrt{\log(1/\delta^2)} \right) \triangleq C_{T,\mu,\delta}. \quad (39)$$

*Proof.* It is straightforward to verify that

$$\begin{aligned} & \mathbb{P} \left( \{x : |x - \mu| \leq \sqrt{\log(1/\delta^2)}\} \cup \{x : |x + \mu| \leq \sqrt{\log(1/\delta^2)}\} \right) \\ &= \mathbb{P}\{x : |x - \mu| \leq \sqrt{\log(1/\delta^2)}\} + \mathbb{P}\{x : |x + \mu| \leq \sqrt{\log(1/\delta^2)}\} \\ &= \int_{\mu - \sqrt{\log(1/\delta^2)}}^{\mu + \sqrt{\log(1/\delta^2)}} p_0(x) dx + \int_{-\mu - \sqrt{\log(1/\delta^2)}}^{-\mu + \sqrt{\log(1/\delta^2)}} p_0(x) dx \\ &\geq q_2 \int_{\mu - \sqrt{\log(1/\delta^2)}}^{\mu + \sqrt{\log(1/\delta^2)}} \mathcal{N}(x; \mu, 1) dx + q_1 \int_{-\mu - \sqrt{\log(1/\delta^2)}}^{-\mu + \sqrt{\log(1/\delta^2)}} \mathcal{N}(x; -\mu, 1) dx \\ &= q_2 \int_{-\sqrt{\log(1/\delta^2)}}^{\sqrt{\log(1/\delta^2)}} \mathcal{N}(x; 0, 1) dx + q_1 \int_{-\sqrt{\log(1/\delta^2)}}^{\sqrt{\log(1/\delta^2)}} \mathcal{N}(x; 0, 1) dx \\ &= (q_1 + q_2) \cdot \mathbb{P}\{\epsilon : |\epsilon| \leq \sqrt{\log(1/\delta^2)}\} \geq 1 - \delta, \end{aligned}$$

where the last inequality applies (21). That is, for any  $\delta > 0$ ,  $\delta \ll 1$ , with the probability of at least  $1 - \delta$ , we have

$$|x| \in [\mu - \sqrt{\log(1/\delta^2)}, \mu + \sqrt{\log(1/\delta^2)}] = \Theta(\mu). \quad (40)$$

Hence, Lemma 1 gives

$$|x(t)| \lesssim C_T \left( \mu + \sqrt{\log(1/\delta^2)} \right), \quad (41)$$

which completes the proof.  $\square$

**Lemma 8** (Monte Carlo estimates, Gaussian mixtures). *Define the Monte Carlo error*

$$\text{Err}_{\text{MC}} = \text{Err}_{\text{MC}}(\boldsymbol{\theta}, \bar{\boldsymbol{\theta}}; T, \lambda(\cdot)) := \mathbb{E}_{t \sim \mathcal{U}(0, T)} \left[ \lambda(t) \cdot \mathbb{E}_{x(t) \sim p_t} \left[ \|s_{t, \boldsymbol{\theta}}(x(t)) - \bar{s}_{t, \bar{\boldsymbol{\theta}}}(x(t))\|_2^2 \right] \right]. \quad (42)$$

*Suppose that the trainable parameter  $\mathbf{a}$  and embedding function  $\mathbf{e}(\cdot)$  are both bounded, and  $x$  is sampled from a one-dimensional 2-mode Gaussian mixture:  $p_0(x) = q_1\mathcal{N}(x; -\mu, 1) + q_2\mathcal{N}(x; \mu, 1)$ , where  $\mu > 0$ ,  $q_1, q_2 > 0$  with  $q_1 + q_2 = 1$  are all constants. Then, given any  $\boldsymbol{\theta}$ , for any  $\delta > 0$ ,  $\delta \ll 1$ ,  $\mu > \sqrt{\log(1/\delta^2)}$ , with the probability of at least  $1 - \delta$ , there exists  $\boldsymbol{\theta}$  such that*

$$\text{Err}_{\text{MC}} \lesssim \mu^2 \frac{\log(1/\delta)}{m}, \quad (43)$$

where  $\lesssim$  hides universal positive constants only depending on  $T$ .

*Proof.* According to Lemma 7, we just need to follow the proof of Lemma 6 by replacing  $C_{T,\delta}$  by  $C_{T,\mu,\delta}$ . Notably, based on (33) in the proof of Lemma 6, one can finally derive

$$\text{Err}_{\text{MC}} \lesssim \mu^2 \frac{\log(1/\delta)}{m},$$

which gives the desired estimates.  $\square$

*Proof of Theorem 2.* We decompose the loss  $\tilde{\mathcal{L}}(\hat{\boldsymbol{\theta}}_n(\tau))$  in the same way as in the proof of Theorem 1. In fact, Theorem 2 is similarly proved by replacing  $C_{T,\delta}$  in the proof of Theorem 1 by  $C_{T,\mu,\delta}$ . Note that  $C_{T,\mu,\delta} = C_T \left( \mu + \sqrt{\log(1/\delta^2)} \right) \lesssim \mu$  for  $\mu \gg 1$ , the proof is completed.  $\square$

**Remark 6.** A standard variance is used here for convenience. In general, if  $\text{var} = o(\mu)$  as  $\mu \rightarrow +\infty$  (e.g. a bounded var), similar analysis and results are supposed to hold. However, this is different when  $\text{var} = \Theta(\mu)$  as  $\mu \rightarrow +\infty$ , since the modes are not separated in this case, and we can not characterize modes shift by simply varying  $\mu$ .

**Remark 7.** Simply scaling down inputs seems not to resolve the adverse effect of modes shift, since the ground truth  $\mu$  is unknown. One can use the input scale to approximate  $\mu$  on toy datasets, but it is not that trivial in practice, particularly for real-world applications with multiple high-dimensional modes in varied scales.

**The model-target inconsistency.** Informally, there is inconsistency between the score network model and target score function. In fact, given the target Gaussian mixture  $p_0(x) = q_1\mathcal{N}(x; -\mu, 1) + q_2\mathcal{N}(x; \mu, 1)$ , the target score function is

$$s_0(x) = -x + \frac{q_2\mathcal{N}(x; \mu, 1) - q_1\mathcal{N}(x; -\mu, 1)}{q_2\mathcal{N}(x; \mu, 1) + q_1\mathcal{N}(x; -\mu, 1)}\mu,$$

which gives  $s_0(x) \approx -x + \mu, x \geq 0$  and  $s_0(x) \approx -x - \mu, x \leq 0$ . While the score network model is

$$s_{0,\theta}(x) \approx \left( \frac{1}{m} \sum_{i:w_i>0} a_i w_i \right) x + \left( \frac{1}{m} \sum_{i:w_i>0} a_i \mathbf{u}_i^\top \mathbf{e}(0) \right),$$

where " $\approx$ " holds since  $|x| = \Theta(\mu)$  by (40) ( $\mu \gg 1$ ). Both  $s_0$  and  $s_{0,\theta}$  are linear functions, but they have unmatched scales in slopes and intercepts:  $O(1)$  and  $O(\mu)$  for  $s_0$ , but both  $O(a)$ 's for  $s_{0,\theta}$ . That is, modeling Gaussian mixtures with large modes distances as random feature models is inconsistent.

**Remark 8.** Recall that the goal of this work is to estimate  $D_{\text{KL}}(p_0 \| p_{0,\hat{\theta}_n(\tau)})$ , which aims at the density estimation characterization. This is different from [32], where the equivalence of learned and target data manifolds (support sets) is studied.

## B Additional Experiments

### B.1 Early-Stopping Generalization

We illustrate the early-stopping generalization gap established in Theorem 1 using the Adam optimizer. All the configurations remain the same as Section 4.1 except that the learning rate is now  $10^{-3}$ .

From Figure 8, one can observe that the KL divergence achieves its minimum at the 1000th training epoch, and it starts to increase after this turning point. The plot is aligned with Theorem 1, which states that there exists an optimal early-stopping time when the model can generalize well, indicating the effectiveness of the upper bound. Further, the KL divergence begins to oscillate after the minimum point (1000th training epoch), which may suggest a phase transition in the KL divergence dynamics, and the transition point is around the optimal early-stopping time. The finding is aligned with SGD setting.

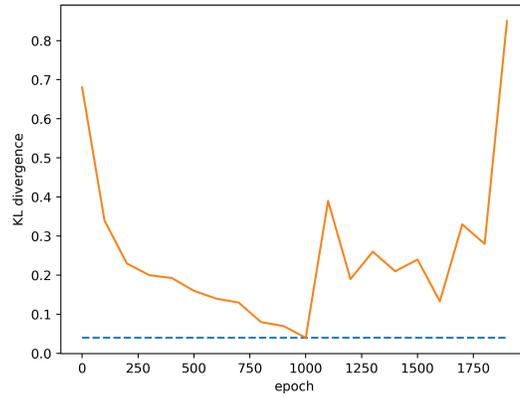


Figure 8: The KL divergence dynamics under the Adam optimizer.

### B.2 Modes Shift Effect

We further test the relationship between the modes' distance and the generalization performance using the Adam optimizer under the same configurations.

In Figure 9 and Figure 10, it is shown that the training of modeled distributions exhibits the same two-stage dynamics as the SGD setting, indicating that *the modes shift effect holds not particularly for a certain optimizer.*

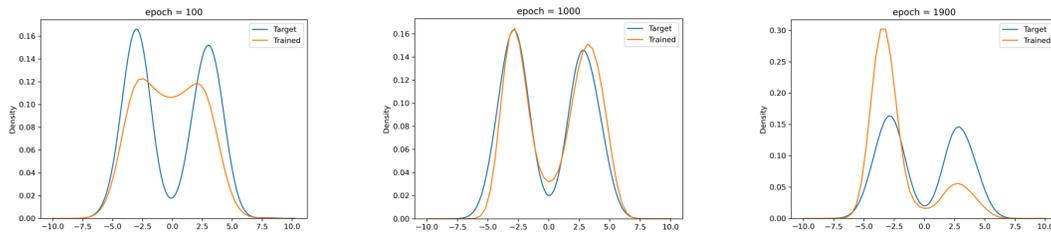


Figure 9: The Adam training dynamics when the distance between two modes is 6 ( $\mu = 3$ ).

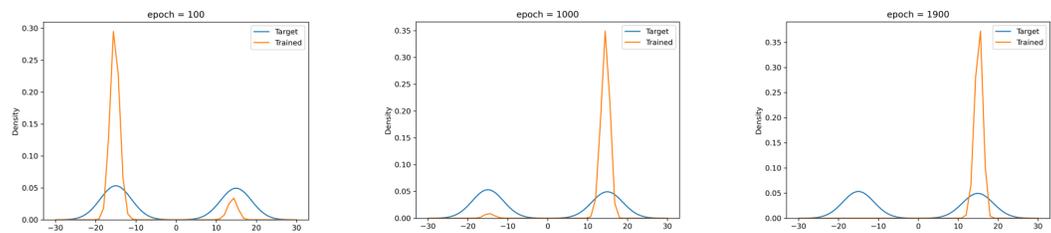


Figure 10: The Adam training dynamics when the distance between two modes is 30 ( $\mu = 15$ ).

### B.3 Model Capacity Dependency

We also numerically study the dependency of generalization on the model capacity. Following the same configurations in Section 4.1.1, we conduct experiments for different hidden dimensions varying from  $2^1$  to  $2^{11}$ .

In Figure 11, the left plot shows the KL divergence from the trained distribution at the 1000th training epoch to the target distribution, and the right plot shows the time duration of the modeled distribution to generalize. Here, the generalization criterion we select is  $D_{\text{KL}} \leq 10^{-1}$ , and we stop training at epoch = 10000. Both the two plots in Figure 11 indicate that increasing model capacity benefits the generalization, which also verifies the corresponding theoretical results ( $m$ -dependency in Theorem 1 and Theorem 2).

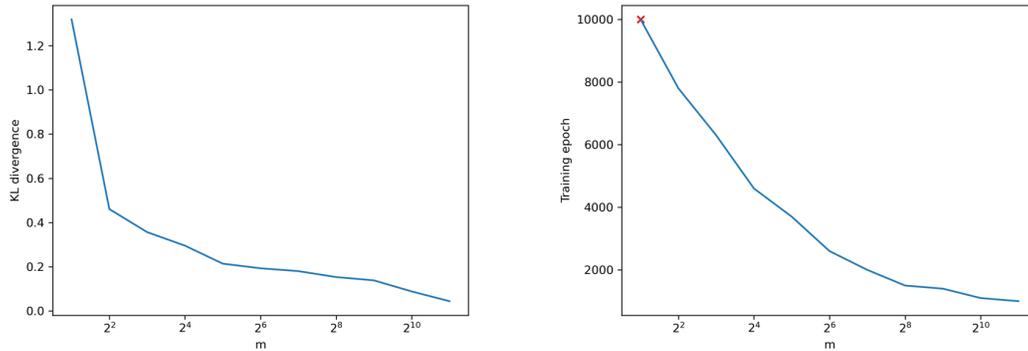


Figure 11: Left: The KL divergence from the model trained after 1000 epochs to the target distribution under different hidden dimensions. Right: The earliest training epoch when  $D_{\text{KL}} \leq 10^{-1}$ . Here, “x” means that the model does not generalize when the training is stopped at the maximum training epoch 10000.