

# A Note on the Convergence of Denoising Diffusion Probabilistic Models

Sokhna Diarra Mbacke  
*Université Laval*

*sokhna-diarra.mbacke.1@ulaval.ca*

Omar Rivasplata  
*University College London*

*o.rivasplata@ucl.ac.uk*

Reviewed on OpenReview: <https://openreview.net/forum?id=wLe1bG93yc>

## Abstract

Diffusion models are one of the most important families of deep generative models. In this note, we derive a quantitative upper bound on the Wasserstein distance between the target distribution and the distribution learned by a diffusion model. Unlike previous works on this topic, our result does not make assumptions on the learned score function. Moreover, our result holds for arbitrary data-generating distributions on bounded instance spaces, even those without a density with respect to Lebesgue measure, and the upper bound does not suffer from exponential dependencies on the ambient space dimension. Our main result builds upon the recent work of Mbacke et al. (2023) and our proofs are elementary.

## 1 Introduction

Along with generative adversarial networks (Goodfellow et al., 2014) and variational autoencoders (VAEs) (Kingma & Welling, 2014; Rezende et al., 2014), diffusion models (Sohl-Dickstein et al., 2015; Song & Ermon, 2019; Ho et al., 2020) are one of the most prominent families of deep generative models. They have exhibited impressive empirical performance in image (Dhariwal & Nichol, 2021; Ho et al., 2022) and audio (Chen et al., 2021; Popov et al., 2021) generation, as well as other applications (Zhou et al., 2021; Sasaki et al., 2021; Li et al., 2022; Trabucco et al., 2023).

There are two main approaches to diffusion models: denoising diffusion probabilistic models (DDPMs) (Sohl-Dickstein et al., 2015; Ho et al., 2020) and score-based generative models (Song & Ermon, 2019) (SGMs). The former kind, DDPMs, progressively transform samples from the target distribution into noise through a forward process, and train a backward process that reverses the transformation and is used to generate new samples. On the other hand, SGMs use score matching techniques (Hyvärinen & Dayan, 2005; Vincent, 2011) to learn an approximation of the score function of the data-generating distribution, then generate new samples using Langevin dynamics. Since for real-world distributions the score function might not exist, Song & Ermon (2019) propose adding different noise levels to the training samples to cover the whole instance space, and train a neural network to simultaneously learn the score function for all noise levels.

Although DDPMs and SGMs might appear to be different approaches at first, Ho et al. (2020) showed that DDPMs implicitly learn an approximation of the score function and the sampling process resembles Langevin dynamics. Furthermore, Song et al. (2021b) derived a unifying view of both techniques using stochastic differential equations (SDEs). The SGM of Song & Ermon (2019) can be seen as a discretization of the Brownian motion process, and the DDPM of Ho et al. (2020) as a discretization of an Ornstein–Uhlenbeck process. Hence, both DDPMs and SGMs are usually referred to as SGMs in the literature. This explains why the previous works studying the theoretical properties of diffusion models utilize the score-based formulation, which requires assumptions on the performance of the learned score function.

In this work, we take a different approach and apply techniques developed by Mbacke et al. (2023) for VAEs to DDPMs, which can be seen as hierarchical VAEs with fixed encoders (Luo, 2022). This approach allows us to derive quantitative Wasserstein-based upper bounds, with no assumptions on the data distribution, no assumptions on the learned score function, and elementary proofs that do not require the SDE toolbox. Moreover, our bounds do not suffer from any costly discretization step, such as the one in De Bortoli (2022), since we consider the forward and backward processes as being discrete-time from the outset, instead of seeing them as discretizations of continuous-time processes.

## 1.1 Related Works

There has been a growing body of work aiming to establish theoretical results on the convergence of SGMs (Block et al., 2020; De Bortoli et al., 2021; Song et al., 2021a; Lee et al., 2022; De Bortoli, 2022; Kwon et al., 2022; Lee et al., 2023; Chen et al., 2023; Li et al., 2023; Benton et al., 2023), but these works either rely on strong assumptions on the data-generating distribution, derive non quantitative upper bounds, or suffer from exponential dependencies on some of the parameters. We manage to avoid all three of these pitfalls. The bounds of Lee et al. (2022) rely on very strong assumptions on the data-generating distribution, such as log-Sobolev inequalities, which are not realistic for real-world data distributions. Furthermore, Song et al. (2021a); Chen et al. (2023); Lee et al. (2023) establish upper bounds on the Kullback-Leibler (KL) divergence or the total variation (TV) distance between the data-generating distribution and the distribution learned by the diffusion model; however, as noted by Pidstrigach (2022) and Chen et al. (2023), unless one makes strong assumptions on the support of the data-generating distribution, KL and TV reach their maximum values. Such assumptions arguably do not hold for real-world data-generating distributions which are widely believed to satisfy the manifold hypothesis (Narayanan & Mitter, 2010; Fefferman et al., 2016; Pope et al., 2021). The work of Pidstrigach (2022) establishes conditions under which the support of the input distribution is equal to the support of the learned distribution, and generalizes the bound of Song et al. (2021a) to all  $f$ -divergences. Assuming  $L_2$  accurate score estimation, Chen et al. (2023) and Lee et al. (2023) establish Wasserstein distance upper bounds under weaker assumptions on the data-generating distribution, but their Wasserstein-based bounds are not quantitative. De Bortoli (2022) derives quantitative Wasserstein distance upper bounds under the manifold hypothesis, but their bounds suffer from exponential dependencies on some of the problem parameters.

## 1.2 Our contributions

In this work, we avoid strong assumptions on the data-generating distribution, and establish a quantitative Wasserstein distance upper bound without exponential dependencies on problem parameters including ambient space dimension. Moreover, a common thread in the works cited above is that their bounds depend on the error of the score estimator. According to Chen et al. (2023), “Providing precise guarantees for estimation of the score function is difficult, as it requires an understanding of the non-convex training dynamics of neural network optimization that is currently out of reach.” Hence, we derive upper bounds without assumptions on the learned score function. Instead, our bound depends on a *reconstruction loss* computed on a finite i.i.d. sample. Intuitively, we define a loss function  $\ell(\mathbf{x}_T, \mathbf{x}_0)$  (see equation 6), which measures the average Euclidean distance between a sample  $\mathbf{x}_0$  from the data-generating distribution, and the reconstruction  $\hat{\mathbf{x}}_0$  obtained by sampling noise  $\mathbf{x}_T \sim q(\mathbf{x}_T/\mathbf{x}_0)$  and passing it through the backward process (parameterized by  $\theta$ ). This approach is motivated by the work of Mbacke et al. (2023) on VAEs.

There are many advantages to this approach: no restrictive assumptions on the data-generating distribution, no exponential dependencies on the dimension, and a quantitative upper bound based on the Wasserstein distance. Moreover, our approach has the benefit of utilizing very simple and elementary proofs.

## 2 Preliminaries

Throughout the paper, we use lower-case letters to denote both probability measures and their densities w.r.t. the Lebesgue measure, and we add variables in parentheses to improve readability (e.g.  $q(\mathbf{x}_t/\mathbf{x}_{t-1})$  to indicate a time-dependent conditional distribution). We consider an instance space  $X$  which is a subset of  $\mathbb{R}^D$  with




Figure 1: Denoising diffusion model

the Euclidean distance as underlying metric, and a target data-generating distribution  $\mu \in \mathcal{M}_+^1(X)$ . Notice that we do not assume  $\mu$  has a density w.r.t. the Lebesgue measure. Moreover,  $\|\cdot\|$  denotes the Euclidean ( $L_2$ ) norm and we write  $\mathbb{E}_{p(\mathbf{x})}$  as a shorthand for  $\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})}$ . Given probability measures  $p, q \in \mathcal{M}_+^1(X)$  and a real number  $k > 1$ , the Wasserstein distance of order  $k$  is defined as (Villani, 2009):

$$W_k(p, q) = \inf_{(\rho, \gamma)} \int_X \|\mathbf{x} - \mathbf{y}\|^k d(\rho, \gamma)^{1/k},$$

where  $\Gamma(p, q)$  denotes the set of couplings of  $p$  and  $q$ , meaning the set of joint distributions on  $X \times X$  with respective marginals  $p$  and  $q$ . We refer to the product measure  $p \otimes q$  as the *trivial coupling*, and we refer to the Wasserstein distance of order 1 simply as *the Wasserstein distance*.

## 2.1 Denoising Diffusion Models

Instead of using the SDE arsenal, we present diffusion models using the DDPM formulation with discrete-time processes. A diffusion model comprises two discrete-time stochastic processes: a forward process and a backward process. Both processes are indexed by time  $0 \leq t \leq T$ , where the number of time-steps  $T$  is a pre-set choice. See Figure 1 for an illustration, and Luo (2022) for a detailed tutorial.

**The forward process.** The forward process transforms a datapoint  $\mathbf{x}_0 \sim \mu$  into a noise distribution  $q(\mathbf{x}_T/\mathbf{x}_0)$ , via a sequence of conditional distributions  $q(\mathbf{x}_t/\mathbf{x}_{t-1})$  for  $1 \leq t \leq T$ . It is assumed that the forward process is defined so that for large enough  $T$ , the distribution  $q(\mathbf{x}_T/\mathbf{x}_0)$  is close to a simple noise distribution  $p(\mathbf{x}_T)$  which is referred to as the *prior distribution*. For instance, Ho et al. (2020) chose  $p(\mathbf{x}_T) = \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$ , the standard multivariate normal distribution.

**The backward process.** The backward process is a Markov process with parametric transition kernels. The goal of the backward process is to implement the reverse action of the forward process: transforming noise samples into (approximate) samples from the distribution  $\mu$ . Following Ho et al. (2020), we assume the backward process to be defined by Gaussian distributions  $p(\mathbf{x}_{t-1}/\mathbf{x}_t)$  defined for  $2 \leq t \leq T$  as

$$p(\mathbf{x}_{t-1}/\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; g^t(\mathbf{x}_t), \frac{2}{t}\mathbf{I}), \quad (1)$$

and

$$p(\mathbf{x}_0/\mathbf{x}_1) = g^1(\mathbf{x}_1), \quad (2)$$

where the variance parameters  $\frac{2}{1}, \dots, \frac{2}{T} \in \mathbb{R}_{>0}$  are defined by a fixed schedule, the mean functions  $g^t : \mathbb{R}^D \rightarrow \mathbb{R}^D$  are learned using a neural network (with parameters  $\theta$ ) for  $2 \leq t \leq T$ , and  $g^1 : \mathbb{R}^D \rightarrow X$  is a separate function dependent on  $\theta_1$ . In practice, Ho et al. (2020) used the same network for the functions  $g^t$  for  $2 \leq t \leq T$ , and a separate discrete decoder for  $g^1$ .

Generating new samples from a trained diffusion model is done by sampling  $\mathbf{x}_{t-1} \sim p(\mathbf{x}_{t-1}/\mathbf{x}_t)$  for  $1 \leq t \leq T$ , starting from a noise vector  $\mathbf{x}_T \sim p(\mathbf{x}_T)$  sampled from the prior  $p(\mathbf{x}_T)$ .

We make the following assumption on the backward process.

**Assumption 1.** We assume for each  $1 \leq t \leq T$  there exists a constant  $K^t > 0$  such that for every  $\mathbf{x}_1, \mathbf{x}_2 \in X$ ,

$$\|g^t(\mathbf{x}_1) - g^t(\mathbf{x}_2)\| \leq K^t \|\mathbf{x}_1 - \mathbf{x}_2\|.$$

In other words,  $g^t$  is  $K^t$ -Lipschitz continuous. We discuss this assumption in Remark 3.2.

## 2.2 Additional Definitions

We define the distribution  $\mathbb{P}(\cdot/\mathbf{x}_0)$  as

$$\mathbb{P}(\cdot/\mathbf{x}_0) = q(\mathbf{x}_T/\mathbf{x}_0)\rho(\mathbf{x}_{T-1}/\mathbf{x}_T)\rho(\mathbf{x}_{T-2}/\mathbf{x}_{T-1})\dots\rho(\mathbf{x}_1/\mathbf{x}_2)\rho(\cdot/\mathbf{x}_1). \quad (3)$$

Intuitively, for each  $\mathbf{x}_0 \in X$ ,  $\mathbb{P}(\cdot/\mathbf{x}_0)$  denotes the distribution on  $X$  obtained by reconstructing samples from  $q(\mathbf{x}_T/\mathbf{x}_0)$  through the backward process. Another way of seeing this distribution is that for any function  $f: X \rightarrow \mathbb{R}$ , the following equation holds:<sup>1</sup>

$$\mathbb{E}_{(\hat{\mathbf{x}}_0/\mathbf{x}_0)} f(\hat{\mathbf{x}}_0) = \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0)} \mathbb{E}_{\rho(\mathbf{x}_{T-1}/\mathbf{x}_T)} \dots \mathbb{E}_{\rho(\mathbf{x}_1/\mathbf{x}_2)} \mathbb{E}_{\rho(\hat{\mathbf{x}}_0/\mathbf{x}_1)} f(\hat{\mathbf{x}}_0). \quad (4)$$

Given a finite set  $S = \{\mathbf{x}_0^1, \dots, \mathbf{x}_0^n\} \stackrel{\text{iid}}{\sim} \mu$ , we define the regenerated distribution as the following mixture:

$$\mu^n = \frac{1}{n} \sum_{i=1}^n \mathbb{P}(\cdot/\mathbf{x}_0^i). \quad (5)$$

This definition is analogous to the empirical regenerated distribution defined by Mbacke et al. (2023) for VAEs. The distribution on  $X$  learned by the diffusion model is denoted  $\mathbb{P}(\cdot)$  and defined as

$$\mathbb{P}(\cdot) = \rho(\mathbf{x}_T)\rho(\mathbf{x}_{T-1}/\mathbf{x}_T)\rho(\mathbf{x}_{T-2}/\mathbf{x}_{T-1})\dots\rho(\mathbf{x}_1/\mathbf{x}_2)\rho(\cdot/\mathbf{x}_1).$$

In other words, for any function  $f: X \rightarrow \mathbb{R}$ , the expectation of  $f$  w.r.t.  $\mathbb{P}(\cdot)$  is

$$\mathbb{E}_{(\hat{\mathbf{x}}_0)} f(\hat{\mathbf{x}}_0) = \mathbb{E}_{\rho(\mathbf{x}_T)} \mathbb{E}_{\rho(\mathbf{x}_{T-1}/\mathbf{x}_T)} \dots \mathbb{E}_{\rho(\mathbf{x}_1/\mathbf{x}_2)} \mathbb{E}_{\rho(\hat{\mathbf{x}}_0/\mathbf{x}_1)} f(\hat{\mathbf{x}}_0).$$

Hence, both  $\mathbb{P}(\cdot)$  and  $\mathbb{P}(\cdot/\mathbf{x}_0)$  are defined using the backward process, with the difference that  $\mathbb{P}(\cdot)$  starts with the prior  $\rho(\mathbf{x}_T) = \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$  while  $\mathbb{P}(\cdot/\mathbf{x}_0)$  starts with the noise distribution  $q(\mathbf{x}_T/\mathbf{x}_0)$ .

Finally, we define the loss function  $\ell: X \times X \rightarrow \mathbb{R}$  as

$$\ell(\mathbf{x}_T, \mathbf{x}_0) = \mathbb{E}_{\rho(\mathbf{x}_{T-1}/\mathbf{x}_T)} \mathbb{E}_{\rho(\mathbf{x}_{T-2}/\mathbf{x}_{T-1})} \dots \mathbb{E}_{\rho(\mathbf{x}_1/\mathbf{x}_2)} \mathbb{E}_{\rho(\hat{\mathbf{x}}_0/\mathbf{x}_1)} \|\mathbf{x}_0 - \hat{\mathbf{x}}_0\|. \quad (6)$$

Hence, given a noise vector  $\mathbf{x}_T$  and a sample  $\mathbf{x}_0$ , the loss  $\ell(\mathbf{x}_T, \mathbf{x}_0)$  denotes the average Euclidean distance between  $\mathbf{x}_0$  and any sample obtained by passing  $\mathbf{x}_T$  through the backward process.

## 2.3 Our Approach

The goal is to upper-bound the distance  $W_1(\mu, \mathbb{P}(\cdot))$ . Since the triangle inequality implies

$$W_1(\mu, \mathbb{P}(\cdot)) \leq W_1(\mu, \mu^n) + W_1(\mu^n, \mathbb{P}(\cdot)), \quad (7)$$

we can upper-bound the distance  $W_1(\mu, \mathbb{P}(\cdot))$  by upper-bounding the two expressions on the right-hand side of equation 7 separately. The upper bound on  $W_1(\mu, \mu^n)$  is obtained using a straightforward adaptation of a proof that was developed by Mbacke et al. (2023). First,  $W_1(\mu, \mu^n)$  is upper-bounded using the expectation of the loss function  $\ell$ , then the resulting expression is upper-bounded using a PAC-Bayesian-style expression dependent on the empirical risk and the prior-matching term.

The upper bound on the second term  $W_1(\mu^n, \mathbb{P}(\cdot))$  uses the definition of  $\mu^n$ . Intuitively, the difference between  $\mathbb{P}(\cdot/\mathbf{x}_0^i)$  and  $\mathbb{P}(\cdot)$  is determined by the corresponding initial distributions:  $q(\mathbf{x}_T/\mathbf{x}_0^i)$  for  $\mathbb{P}(\cdot/\mathbf{x}_0^i)$ , and  $\rho(\mathbf{x}_T)$  for  $\mathbb{P}(\cdot)$ . Hence, if the two initial distributions are close, and if the steps of the backward process are smooth (see Assumption 1), then  $\mathbb{P}(\cdot/\mathbf{x}_0^i)$  and  $\mathbb{P}(\cdot)$  are close to each other.

<sup>1</sup>More formally, we give a definition of  $\pi_\theta(\cdot/\mathbf{x}_0)$  via expectations of test functions by requiring that equation 4 holds for every function  $f$  in some appropriate measure-determining function class.

### 3 Main Result

#### 3.1 Theorem Statement

We are now ready to state our main result: a quantitative upper bound on the Wasserstein distance between the data-generating distribution  $\mu$  and the learned distribution  $\hat{\mu}$  ( $\cdot$ ).

**Theorem 3.1.** *Assume the instance space  $X$  has finite diameter  $\Delta = \sup_{\mathbf{x}, \mathbf{x}'} \|\mathbf{x} - \mathbf{x}'\| < \infty$ , and let  $\eta > 0$  and  $\epsilon \in (0, 1)$  be real numbers. Using the definitions and assumptions of the previous section, the following inequality holds with probability at least  $1 - \epsilon$  over the random draw of  $S = \{\mathbf{x}_0^1, \dots, \mathbf{x}_0^n\} \stackrel{iid}{\sim} \mu$ :*

$$W_1(\mu, \hat{\mu}) \leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0^i)} \|\mathbf{x}_T, \mathbf{x}_0^i\| + \frac{1}{n} \sum_{i=1}^n \text{KL}(q(\mathbf{x}_T/\mathbf{x}_0^i) \parallel p(\mathbf{x}_T)) + \log \frac{1}{\epsilon} + \frac{\Delta^2}{8n} + \sum_{t=1}^T K^t \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0^i)} \mathbb{E}_{p(\mathbf{y}_T)} \|\mathbf{x}_T - \mathbf{y}_T\| + \sum_{t=2}^T K^t \sum_{i=1}^n \mathbb{E}_{\epsilon, \epsilon'} \|\epsilon - \epsilon'\|. \quad (8)$$

Where  $\epsilon, \epsilon' \sim N(\mathbf{0}, \mathbf{I})$  are standard Gaussian vectors.

**Remark 3.1.** Before presenting the proof, let us discuss Theorem 3.1.

- Because the right-hand side of equation 8 depends on a quantity computed using a finite i.i.d. sample  $S$ , the bound holds with high probability w.r.t. the randomness of  $S$ . This is the price we pay for having a quantitative upper bound with no exponential dependencies on problem parameters and no assumptions on the data-generating distribution  $\mu$ .
- The first term of the right-hand side of equation 8 is the average reconstruction loss computed over the sample  $S = \{\mathbf{x}_0^1, \dots, \mathbf{x}_0^n\}$ . Note that for each  $1 \leq i \leq n$ , the expectation of  $\|\mathbf{x}_T/\mathbf{x}_0^i\|$  is only computed w.r.t. the noise distribution  $q(\mathbf{x}_T/\mathbf{x}_0^i)$  defined by  $\mathbf{x}_0^i$  itself. Hence, this term measures how well a noise vector  $\mathbf{x}_T \sim q(\mathbf{x}_T/\mathbf{x}_0^i)$  recovers the original sample  $\mathbf{x}_0^i$  using the backward process, and averages over the set  $S = \{\mathbf{x}_0^1, \dots, \mathbf{x}_0^n\}$ .
- If the Lipschitz constants satisfy  $K^t < 1$  for all  $1 \leq t \leq T$ , then the larger  $T$  is, the smaller the upper bound gets. This is because the product of  $K^t$ 's then converges to 0. In Remark 3.2 below, we show that the assumption that  $K^t < 1$  for all  $t$  is a quite reasonable one.
- The hyperparameter  $\epsilon$  controls the trade-off between the prior-matching (KL) term and the diameter term  $\frac{\Delta^2}{8n}$ . If  $K^t < 1$  for all  $1 \leq t \leq T$  and  $T \rightarrow \infty$ , then the convergence of the bound largely depends on the choice of  $\epsilon$ . In that case,  $\epsilon = n^{-1/2}$  leads to a faster convergence, while  $\epsilon = n^{-1}$  leads to a slower convergence to a smaller quantity. This is because the bound of Mbacke et al. (2023) stems from PAC-Bayesian theory, where this trade-off is common, see e.g. Alquier (2021).
- The last term of equation 8 does not depend on the sample size  $n$ . Hence, the upper bound given by Theorem 3.1 does not converge to 0 as  $n \rightarrow \infty$ . However, if the Lipschitz factors  $(K^t)_{1 \leq t \leq T}$  are all less than 1, then this term can be very small, specially in low dimensional spaces.

#### 3.2 Proof of the main theorem

The following result is an adaptation of a result by Mbacke et al. (2023).

**Lemma 3.2.** *Let  $\eta > 0$  and  $\epsilon \in (0, 1)$  be real numbers. With probability at least  $1 - \epsilon$  over the randomness of the sample  $S = \{\mathbf{x}_0^1, \dots, \mathbf{x}_0^n\} \stackrel{iid}{\sim} \mu$ , the following holds:*

$$W_1(\mu, \mu^n) \leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0^i)} \|\mathbf{x}_T, \mathbf{x}_0^i\| + \frac{1}{n} \sum_{i=1}^n \text{KL}(q(\mathbf{x}_T/\mathbf{x}_0^i) \parallel p(\mathbf{x}_T)) + \log \frac{1}{\epsilon} + \frac{\Delta^2}{8n}. \quad (9)$$

The proof of this result is a straightforward adaptation of Mbacke et al. (2023, Lemma D.1). We provide the proof in the supplementary material (Section A.1) for completeness.

Now, let us focus our attention on the second term of the right-hand side of equation 7, namely  $W_1(\mu^n, (\cdot))$ . This part is trickier than for VAEs, for which the generative model's distribution is simply a pushforward measure. Here, we have a non-deterministic sampling process with  $T$  steps.

Assumption 1 leads to the following lemma on the backward process.

**Lemma 3.3.** *For any given  $\mathbf{x}_1, \mathbf{y}_1 \sim \mathcal{X}$  we have*

$$\mathbb{E}_{p(\mathbf{x}_0/\mathbf{x}_1)} \mathbb{E}_{p(\mathbf{y}_0/\mathbf{y}_1)} \|\mathbf{x}_0 - \mathbf{y}_0\| \leq K^1 \|\mathbf{x}_1 - \mathbf{y}_1\|.$$

Moreover, if  $2 \leq t \leq T$ , then for any given  $\mathbf{x}_t, \mathbf{y}_t \sim \mathcal{X}$  we have

$$\mathbb{E}_{p(\mathbf{x}_{t-1}/\mathbf{x}_t)} \mathbb{E}_{p(\mathbf{y}_{t-1}/\mathbf{y}_t)} \|\mathbf{x}_{t-1} - \mathbf{y}_{t-1}\| \leq K^t \|\mathbf{x}_t - \mathbf{y}_t\| + \sum_{\epsilon_t, \epsilon_t} \|\epsilon_t - \epsilon_t\|,$$

where  $\epsilon_t, \epsilon_t \sim N(\mathbf{0}, \mathbf{I})$ , meaning  $\mathbb{E}_{\epsilon_t, \epsilon_t}$  is a shorthand for  $\mathbb{E}_{\epsilon_t, \epsilon_t \sim N(\mathbf{0}, \mathbf{I})}$ .

*Proof.* For the first part, let  $\mathbf{x}_1, \mathbf{y}_1 \sim \mathcal{X}$ . Since according to equation 2 we have  $p(\mathbf{x}_0/\mathbf{x}_1) = g^1(\mathbf{x}_1)(\mathbf{x}_0)$  and  $p(\mathbf{y}_0/\mathbf{y}_1) = g^1(\mathbf{y}_1)(\mathbf{y}_0)$ , then

$$\mathbb{E}_{p(\mathbf{x}_0/\mathbf{x}_1)} \mathbb{E}_{p(\mathbf{y}_0/\mathbf{y}_1)} \|\mathbf{x}_0 - \mathbf{y}_0\| = \|g^1(\mathbf{x}_1) - g^1(\mathbf{y}_1)\| \leq K^1 \|\mathbf{x}_1 - \mathbf{y}_1\|.$$

For the second part, let  $2 \leq t \leq T$  and  $\mathbf{x}_t, \mathbf{y}_t \sim \mathcal{X}$ . Since  $p(\mathbf{x}_{t-1}/\mathbf{x}_t) = N(\mathbf{x}_{t-1}; g^t(\mathbf{x}_t), \frac{1}{2}\mathbf{I})$  by equation 1, the reparameterization trick (Kingma & Welling, 2014) implies that sampling

$$\mathbf{x}_{t-1} \sim p(\mathbf{x}_{t-1}/\mathbf{x}_t)$$

is equivalent to setting

$$\mathbf{x}_{t-1} = g^t(\mathbf{x}_t) + \sum_{\epsilon_t} \epsilon_t, \quad \text{with } \epsilon_t \sim N(\mathbf{0}, \mathbf{I}). \quad (10)$$

Using equation 10, the triangle inequality, and Assumption 1, we obtain

$$\begin{aligned} \mathbb{E}_{p(\mathbf{x}_{t-1}/\mathbf{x}_t)} \mathbb{E}_{p(\mathbf{y}_{t-1}/\mathbf{y}_t)} \|\mathbf{x}_{t-1} - \mathbf{y}_{t-1}\| &= \mathbb{E}_{\epsilon_t} \mathbb{E}_{\epsilon_t} \|g^t(\mathbf{x}_t) + \sum_{\epsilon_t} \epsilon_t - g^t(\mathbf{y}_t) - \sum_{\epsilon_t} \epsilon_t\| \\ &= \mathbb{E}_{\epsilon_t} \mathbb{E}_{\epsilon_t} \|g^t(\mathbf{x}_t) - g^t(\mathbf{y}_t) + \sum_{\epsilon_t} \epsilon_t - \sum_{\epsilon_t} \epsilon_t\| \\ &= \|g^t(\mathbf{x}_t) - g^t(\mathbf{y}_t)\| + \sum_{\epsilon_t} \mathbb{E}_{\epsilon_t} \|\epsilon_t - \epsilon_t\| \\ &\leq K^t \|\mathbf{x}_t - \mathbf{y}_t\| + \sum_{\epsilon_t} \mathbb{E}_{\epsilon_t} \|\epsilon_t - \epsilon_t\|, \end{aligned}$$

where  $\epsilon_t, \epsilon_t \sim N(\mathbf{0}, \mathbf{I})$ . □

Next, we can use the inequalities of Lemma 3.3 to prove the following result.

**Lemma 3.4.** *Let  $T \geq 1$ . The following inequality holds:*

$$\begin{aligned} \mathbb{E}_{p(\mathbf{x}_{T-1}/\mathbf{x}_T)} \mathbb{E}_{p(\mathbf{y}_{T-1}/\mathbf{y}_T)} \mathbb{E}_{p(\mathbf{x}_{T-2}/\mathbf{x}_{T-1})} \mathbb{E}_{p(\mathbf{y}_{T-2}/\mathbf{y}_{T-1})} \dots \mathbb{E}_{p(\mathbf{x}_0/\mathbf{x}_1)} \mathbb{E}_{p(\mathbf{y}_0/\mathbf{y}_1)} \|\mathbf{x}_0 - \mathbf{y}_0\| \\ \leq \sum_{t=1}^T K^t \|\mathbf{x}_T - \mathbf{y}_T\| + \sum_{t=2}^T \sum_{i=1}^{t-1} K^i \sum_{\epsilon_t} \|\epsilon_t - \epsilon_t\|, \end{aligned}$$

where  $\epsilon_t, \epsilon_t \sim N(\mathbf{0}, \mathbf{I})$ .

*Proof Idea.* Lemma 3.4 is proven by induction using Lemma 3.3 in the induction step. The details are in the supplementary material (Section A.2). □

Using the two previous lemmas, we obtain the following upper bound on  $W_1(\mu^n, \cdot)$ .

**Lemma 3.5.** *The following inequality holds:*

$$W_1(\mu^n, \cdot) \leq \sum_{t=1}^T K^t \frac{1}{n} \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0^i)} \mathbb{E}_{p(\mathbf{y}_T)} \|\mathbf{x}_T - \mathbf{y}_T\| + \sum_{t=2}^T \sum_{i=1}^{t-1} K^i \mathbb{E}_{\epsilon, \epsilon} \|\epsilon - \epsilon\|,$$

where  $\epsilon, \epsilon \sim N(\mathbf{0}, \mathbf{I})$ .

*Proof.* Using the definition of  $W_1$ , the trivial coupling, the definitions of  $\mu^n$  and  $\cdot$ , and Lemma 3.4, we get

$$\begin{aligned} W_1(\mu^n, \cdot) &= \inf_{(\mu^n, \cdot)} \int d(\mathbf{x}, \mathbf{y}) d(\mathbf{x}, \mathbf{y}) \\ &= \mathbb{E}_{\mathbf{x} \sim \mu^n} \mathbb{E}_{\mathbf{y}} [\|\mathbf{x} - \mathbf{y}\|] \\ &= \frac{1}{n} \mathbb{E}_{i=1}^n \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0^i)} \mathbb{E}_{p(\mathbf{y}_T)} \mathbb{E}_{p(\mathbf{x}_{T-1}/\mathbf{x}_T)} \mathbb{E}_{p(\mathbf{y}_{T-1}/\mathbf{y}_T)} \dots \mathbb{E}_{p(\mathbf{x}_0/\mathbf{x}_1)} \mathbb{E}_{p(\mathbf{y}_0/\mathbf{y}_1)} \|\mathbf{x}_0 - \mathbf{y}_0\| \\ &= \frac{1}{n} \mathbb{E}_{i=1}^n \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0^i)} \mathbb{E}_{p(\mathbf{y}_T)} \sum_{t=1}^T K^t \|\mathbf{x}_T - \mathbf{y}_T\| + \sum_{t=2}^T \sum_{j=1}^{t-1} K^j \mathbb{E}_{\epsilon, \epsilon} [\|\epsilon - \epsilon\|] \\ &= \sum_{t=1}^T K^t \frac{1}{n} \mathbb{E}_{i=1}^n \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0^i)} \mathbb{E}_{p(\mathbf{y}_T)} \|\mathbf{x}_T - \mathbf{y}_T\| + \sum_{t=2}^T \sum_{i=1}^{t-1} K^i \mathbb{E}_{\epsilon, \epsilon} \|\epsilon - \epsilon\|. \end{aligned}$$

□

Combining Lemmas 3.2 and 3.5 with equation 7 yields Theorem 3.1.

### 3.3 Special case using the forward process of Ho et al. (2020)

Theorem 3.1 establishes a general upper bound that holds for any forward process, as long as the backward process satisfies Assumption 1. In this section, we specialize the statement of the theorem to the particular case of the forward process defined by Ho et al. (2020).

Let  $X \in \mathbb{R}^D$ . In Ho et al. (2020), the forward process is a Gauss-Markov process with transition densities defined as

$$q(\mathbf{x}_t/\mathbf{x}_{t-1}) = N(\mathbf{x}_t; \overline{\mathbf{x}}_{t-1}, (1 - \alpha_t)\mathbf{I}),$$

where  $\alpha_1, \dots, \alpha_T$  is a fixed noise schedule such that  $0 < \alpha_t < 1$  for all  $t$ . This definition implies that at each time step  $1 \leq t \leq T$ ,

$$q(\mathbf{x}_t/\mathbf{x}_0) = N(\mathbf{x}_t; \overline{\mathbf{x}}_t, (1 - \overline{\alpha}_t)\mathbf{I}), \quad \text{with} \quad \overline{\alpha}_t = \prod_{i=1}^t \alpha_i.$$

The optimization objective to train the backward process ensures that for each time step  $t$  the distribution  $p(\mathbf{x}_{t-1}/\mathbf{x}_t)$  remains close to the ground-truth distribution  $q(\mathbf{x}_{t-1}/\mathbf{x}_t, \mathbf{x}_0)$  given by

$$q(\mathbf{x}_{t-1}/\mathbf{x}_t, \mathbf{x}_0) = N(\mathbf{x}_{t-1}; \mu_q^t(\mathbf{x}_t, \mathbf{x}_0), \frac{\alpha_t}{1 - \alpha_t}\mathbf{I}),$$

where

$$\mu_q^t(\mathbf{x}_t, \mathbf{x}_0) = \frac{\overline{\alpha}_t(1 - \overline{\alpha}_{t-1})}{1 - \overline{\alpha}_t} \mathbf{x}_t + \frac{\overline{\alpha}_{t-1}(1 - \alpha_t)}{1 - \overline{\alpha}_t} \mathbf{x}_0. \quad (11)$$

Now, we discuss Assumption 1 under these definitions.

**Remark 3.2.** We can get a glimpse at the range of  $K^t$  for a trained DDPM by looking at the distribution  $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ , since  $p(\mathbf{x}_{t-1}|\mathbf{x}_t)$  is optimized to be as close as possible to  $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ .

For a given  $\mathbf{x}_0$ , let us take a look at the Lipschitz norm of  $\mu_q^t(\mathbf{x}, \mathbf{x}_0)$ . Using equation 11, we have

$$\mu_q^t(\mathbf{x}_t, \mathbf{x}_0) - \mu_q^t(\mathbf{y}_t, \mathbf{x}_0) = \frac{\bar{\sigma}_t(1 - \bar{\sigma}_{t-1})}{1 - \bar{\sigma}_t}(\mathbf{x}_t - \mathbf{y}_t).$$

Hence,  $\mu_q^t(\mathbf{x}, \mathbf{x}_0)$  is  $K_t$ -Lipschitz continuous with

$$K_t = \frac{\bar{\sigma}_t(1 - \bar{\sigma}_{t-1})}{1 - \bar{\sigma}_t}.$$

Now, if  $\bar{\sigma}_t < 1$  for all  $1 \leq t \leq T$ , then we have  $1 - \bar{\sigma}_t > 1 - \bar{\sigma}_{t-1}$  which implies  $K_t < 1$  for all  $1 \leq t \leq T$ .

Remark 3.2 shows that the Lipschitz norm of the mean function  $\mu_q^t(\cdot, \mathbf{x}_0)$  does not depend on  $\mathbf{x}_0$ . Indeed, looking at the previous equation, we can see that for any initial  $\mathbf{x}_0$ , the Lipschitz norm  $K_t = \frac{\bar{\sigma}_t(1 - \bar{\sigma}_{t-1})}{1 - \bar{\sigma}_t}$  only depends on the noise schedule, not  $\mathbf{x}_0$  itself. Since  $g^t$  is optimized to match to  $\mu_q^t(\cdot, \mathbf{x}_0)$ , for each  $\mathbf{x}_0$  in the training set, and all the functions  $\mu_q^t(\cdot, \mathbf{x}_0)$  have the same Lipschitz norm  $K_t$ , we believe it is reasonable to assume  $g^t$  is Lipschitz continuous as well. This is the intuition behind Assumption 1.

**The prior-matching term.** With the definitions of this section, the prior matching term  $\text{KL}(q(\mathbf{x}_T|\mathbf{x}_0) \parallel p(\mathbf{x}_T))$  has the following closed form:

$$\text{KL}(q(\mathbf{x}_T|\mathbf{x}_0) \parallel p(\mathbf{x}_T)) = \frac{1}{2} \left( -D \log(1 - \bar{\sigma}_T) - D^{-T} + \bar{\sigma}_T \|\mathbf{x}_0\|^2 \right).$$

**Upper-bounds on the average distance between Gaussian vectors.** If  $\epsilon, \epsilon'$  are  $D$ -dimensional vectors sampled from  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ , then

$$\mathbb{E}_{\epsilon, \epsilon'} \|\epsilon - \epsilon'\|^2 = 2D.$$

Moreover, since  $q(\mathbf{x}_T|\mathbf{x}_0^i) = \mathcal{N}(\mathbf{x}_T; \bar{\sigma}_T \mathbf{x}_0^i, (1 - \bar{\sigma}_T)\mathbf{I})$  and the prior  $p(\mathbf{y}_T) = \mathcal{N}(\mathbf{y}_T; \mathbf{0}, \mathbf{I})$ ,

$$\mathbb{E}_{q(\mathbf{x}_T|\mathbf{x}_0^i)} \mathbb{E}_{p(\mathbf{y}_T)} \|\mathbf{x}_T - \mathbf{y}_T\|^2 = \bar{\sigma}_T \|\mathbf{x}_0^i\|^2 + (2 - \bar{\sigma}_T)D.$$

**Special case of the main theorem.** With the definitions of this section, the inequality of Theorem 3.1 implies that with probability at least  $1 - \delta$  over the randomness of  $\{\mathbf{x}_0^1, \dots, \mathbf{x}_0^n\}$  iid  $\mu$ :

$$\begin{aligned} \mathcal{W}_1(\mu, (\cdot)) &\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{q(\mathbf{x}_T|\mathbf{x}_0^i)} \|\mathbf{x}_T, \mathbf{x}_0^i\|^2 + \frac{1}{2} \sum_{i=1}^n \left( -D \log(1 - \bar{\sigma}_T) - D^{-T} + \bar{\sigma}_T \|\mathbf{x}_0^i\|^2 + \log \frac{1}{\delta} \right) \\ &\quad + \frac{\Delta^2}{8n} + \sum_{t=1}^T K^t \frac{1}{n} \sum_{i=1}^n \frac{\bar{\sigma}_t(1 - \bar{\sigma}_{t-1})}{1 - \bar{\sigma}_t} \|\mathbf{x}_0^i\|^2 + \frac{\bar{\sigma}_T}{2D} \sum_{t=2}^T \sum_{i=1}^{t-1} K^i \|\mathbf{x}_0^i\|^2. \end{aligned}$$

## 4 Conclusion

This note presents a novel upper bound on the Wasserstein distance between the data-generating distribution and the distribution learned by a diffusion model. Unlike previous works in the field, our main result simultaneously avoids strong assumptions on the data-generating distribution, assumptions on the learned score function, and exponential dependencies, while still providing a quantitative upper bound. However, our bound holds with high probability on the randomness of a finite i.i.d. sample, on which a loss function is computed. Since the loss is a chain of expectations w.r.t. Gaussian distributions, it can either be estimated with high precision or upper bounded using the properties of Gaussian distributions.



## Acknowledgments

This research is supported by the Canada CIFAR AI Chair Program, and the NSERC Discovery grant RGPIN-2020- 07223. The authors sincerely thank Pascal Germain for interesting discussions and suggestions. The first author thanks Mathieu Bazinet and Florence Clerc for proof-reading the manuscript.

## References

- Pierre Alquier. User-friendly introduction to PAC-Bayes bounds. *arXiv preprint arXiv:2110.11216*, 2021.
- Joe Benton, Valentin De Bortoli, Arnaud Doucet, and George Deligiannidis. Nearly  $d$ -Linear Convergence Bounds for Diffusion Models via Stochastic Localization. *arXiv preprint arXiv:2308.03686*, 2023.
- Adam Block, Youssef Mroueh, and Alexander Rakhlin. Generative Modeling with Denoising Auto-Encoders and Langevin Sampling. *arXiv preprint arXiv:2002.00107*, 2020.
- Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan. WaveGrad: Estimating Gradients for Waveform Generation. In *International Conference on Learning Representations*, 2021.
- 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 *The Eleventh International Conference on Learning Representations*, 2023.
- Valentin De Bortoli. Convergence of denoising diffusion models under the manifold hypothesis. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856.
- Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger Bridge with Applications to Score-Based Generative Modeling. In *Advances in Neural Information Processing Systems*, volume 34, pp. 17695–17709, 2021.
- Prafulla Dhariwal and Alexander Nichol. Diffusion Models Beat GANs on Image Synthesis. In *Advances in Neural Information Processing Systems*, volume 34, pp. 8780–8794, 2021.
- Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the Manifold Hypothesis. *Journal of the American Mathematical Society*, 29(4):983–1049, 2016.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks. In *Advances in Neural Information Processing Systems*, volume 27, 2014.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020.
- Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded Diffusion Models for High Fidelity Image Generation. *Journal of Machine Learning Research*, 23(1):2249–2281, 2022.
- Aapo Hyvärinen and Peter Dayan. Estimation of Non-Normalized Statistical Models by Score Matching. *Journal of Machine Learning Research*, 6(4), 2005.
- Diederik P. Kingma and M. Welling. Auto-Encoding Variational Bayes. *CoRR*, abs/1312.6114, 2014.
- Dohyun Kwon, Ying Fan, and Kangwook Lee. Score-based Generative Modeling Secretly Minimizes the Wasserstein Distance. *Advances in Neural Information Processing Systems*, 35:20205–20217, 2022.
- Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence for score-based generative modeling with polynomial complexity. In *Advances in Neural Information Processing Systems*, volume 35, pp. 22870–22882, 2022.
- Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence of score-based generative modeling for general data distributions. In *International Conference on Algorithmic Learning Theory*, pp. 946–985. PMLR, 2023.

- Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. Towards Faster Non-Asymptotic Convergence for Diffusion-Based Generative Models. *arXiv preprint arXiv:2306.09251*, 2023.
- 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.
- Calvin Luo. Understanding Diffusion Models: A Unified Perspective. *arXiv preprint arXiv:2208.11970*, 2022.
- Sokhna Diarra Mbacke, Florence Clerc, and Pascal Germain. Statistical Guarantees for Variational Autoencoders using PAC-Bayesian Theory. In *Advances in Neural Information Processing Systems*, 2023.
- Hariharan Narayanan and Sanjoy Mitter. Sample Complexity of Testing the Manifold Hypothesis. In *Advances in Neural Information Processing Systems*, volume 23, 2010.
- Jakiw Pidstrigach. Score-Based Generative Models Detect Manifolds. In *Advances in Neural Information Processing Systems*, volume 35, pp. 35852–35865, 2022.
- Phil Pope, Chen Zhu, Ahmed Abdelkader, Micah Goldblum, and Tom Goldstein. The Intrinsic Dimension of Images and Its Impact on Learning. In *International Conference on Learning Representations*, 2021.
- Vadim Popov, Ivan Vovk, Vladimir Gogoryan, Tasnima Sadekova, and Mikhail Kudinov. Grad-TTS: A Diffusion Probabilistic Model for Text-to-Speech. In *International Conference on Machine Learning*, pp. 8599–8608. PMLR, 2021.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *International Conference on Machine Learning*, pp. 1278–1286. PMLR, 2014.
- Hiroshi Sasaki, Chris G Willcocks, and Toby P Breckon. UNIT-DDPM: UNpaired Image Translation with Denoising Diffusion Probabilistic Models. *arXiv preprint arXiv:2104.05358*, 2021.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *International Conference on Machine Learning*, pp. 2256–2265. PMLR, 2015.
- Yang Song and Stefano Ermon. Generative Modeling by Estimating Gradients of the Data Distribution. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum Likelihood Training of Score-Based Diffusion Models. In *Advances in Neural Information Processing Systems*, volume 34, 2021a.
- 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*, 2021b.
- Brandon Trabucco, Kyle Doherty, Max Gurinas, and Ruslan Salakhutdinov. Effective Data Augmentation With Diffusion Models. *arXiv preprint arXiv:2302.07944*, 2023.
- Cédric Villani. *Optimal Transport: Old and New*, volume 338. Springer, 2009.
- Pascal Vincent. A Connection Between Score Matching and Denoising Autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.
- Linqi Zhou, Yilun Du, and Jiajun Wu. 3D Shape Generation and Completion through Point-Voxel Diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5826–5835, 2021.

## A Omitted Proofs

### A.1 Proof of Lemma 3.2

Recall Lemma 3.2 states that the following inequality holds with probability  $1 - \delta$ :

$$W_1(\mu, \mu^n) \leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0^i)} \ell(\mathbf{x}_T, \mathbf{x}_0^i) + \frac{1}{n} \sum_{i=1}^n \text{KL}(q(\mathbf{x}_T/\mathbf{x}_0^i) \parallel p(\mathbf{x}_T)) + \log \frac{1}{1 - \delta} + \frac{\Delta^2}{8n}.$$

*Proof of Lemma 3.2.* Using the trivial coupling (product of marginals), the definition of  $\mu^n$  (equation 5), and the definition of the loss function  $\ell$ , we get

$$\begin{aligned} W_1(\mu, \mu^n) &= \inf_{(\mu, \mu^n)} \int_{\mathcal{X} \times \mathcal{X}} \ell(\mathbf{x}, \mathbf{y}) d(\mu \times \mu^n) \\ &= \int_{\mathcal{X}} \int_{\mathcal{X}} \ell(\mathbf{x}, \mathbf{y}) d\mu(\mathbf{x}) d\mu^n(\mathbf{y}) \\ &= \int_{\mathcal{X}} \int_{\mathcal{Y}} \ell(\mathbf{x}, \mathbf{y}) \mu(\mathbf{x}) \mu^n(\mathbf{y}) d\mathbf{y} \\ &= \int_{\mathcal{X}} \mu(\mathbf{x}) \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0^i)} \mathbb{E}_{p(\mathbf{x}_{T-1}/\mathbf{x}_T)} \dots \mathbb{E}_{p(\mathbf{x}_1/\mathbf{x}_2)} \mathbb{E}_{p(\mathbf{y}/\mathbf{x}_1)} \ell(\mathbf{x}, \mathbf{y}) d\mathbf{x} \\ &= \int_{\mathcal{X}} \mu(\mathbf{x}) \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0^i)} \ell(\mathbf{x}_T, \mathbf{x}) d\mathbf{x}. \end{aligned}$$

Using Mbacke et al. (2023, Lemma B.1), the following inequality holds with probability  $1 - \delta$ :

$$\begin{aligned} W_1(\mu, \mu^n) &\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{q(\mathbf{x}_T/\mathbf{x}_0^i)} \ell(\mathbf{x}_T, \mathbf{x}_0^i) + \frac{1}{n} \sum_{i=1}^n \text{KL}(q(\mathbf{x}_T/\mathbf{x}_0^i) \parallel p(\mathbf{x}_T)) + \log \frac{1}{1 - \delta} \\ &\quad + n \log \mathbb{E}_{\mathbf{x}_0} \int_{\mathcal{X}} \int_{\mathcal{X}} \ell(\mathbf{x}, \mathbf{x}') \mu(\mathbf{x}) \mu(\mathbf{x}') d\mathbf{x} d\mathbf{x}' e^{\frac{1}{n} \sum_{i=1}^n \ell(\mathbf{x}_T, \mathbf{x}_0^i)}. \end{aligned} \tag{12}$$

Now, it remains to upper-bound the exponential moment of equation 9. If  $\sup_{\mathbf{x}, \mathbf{x}'} \ell(\mathbf{x}, \mathbf{x}') = \Delta < \infty$ , and  $\delta > 0$  is a real number, then the definition of the loss function  $\ell$  and Hoeffding's lemma yield

$$n \log \mathbb{E}_{\mathbf{x}_0} \int_{\mathcal{X}} \int_{\mathcal{X}} \ell(\mathbf{x}, \mathbf{x}') \mu(\mathbf{x}) \mu(\mathbf{x}') d\mathbf{x} d\mathbf{x}' e^{\frac{1}{n} \sum_{i=1}^n \ell(\mathbf{x}_T, \mathbf{x}_0^i)} \leq n \log \exp \frac{2\Delta^2}{8n^2} = \frac{2\Delta^2}{8n}.$$

□

### A.2 Proof of Lemma 3.4

Recall Lemma 3.4 states that the following inequality holds:

$$\begin{aligned} &\mathbb{E}_{p(\mathbf{x}_{T-1}/\mathbf{x}_T)} \mathbb{E}_{p(\mathbf{y}_{T-1}/\mathbf{y}_T)} \mathbb{E}_{p(\mathbf{x}_{T-2}/\mathbf{x}_{T-1})} \mathbb{E}_{p(\mathbf{y}_{T-2}/\mathbf{y}_{T-1})} \dots \mathbb{E}_{p(\mathbf{x}_0/\mathbf{x}_1)} \mathbb{E}_{p(\mathbf{y}_0/\mathbf{y}_1)} \|\mathbf{x}_0 - \mathbf{y}_0\| \\ &\leq \sum_{t=1}^T K^t \|\mathbf{x}_T - \mathbf{y}_T\| + \sum_{t=2}^T K^t \sum_{i=1}^{t-1} \mathbb{E}_{\epsilon, \epsilon'} \|\epsilon - \epsilon'\|, \end{aligned}$$

*Proof of Lemma 3.4.* Let's do a proof by induction on  $T$ .

- **Base case.**  $T = 1$ . The inequality becomes

$$\mathbb{E}_{\rho(\mathbf{x}_0/\mathbf{x}_1)} \mathbb{E}_{\rho(\mathbf{y}_0/\mathbf{y}_1)} \|\mathbf{x}_0 - \mathbf{y}_0\| \leq K^1 \|\mathbf{x}_1 - \mathbf{y}_1\| \quad (13)$$

which holds by the first part of Lemma 3.3.

- **Induction step.** Assume  $T > 1$ . The induction hypothesis is

$$\mathbb{E}_{\rho(\mathbf{x}_{T-2}/\mathbf{x}_{T-1})} \mathbb{E}_{\rho(\mathbf{y}_{T-2}/\mathbf{y}_{T-1})} \dots \mathbb{E}_{\rho(\mathbf{x}_0/\mathbf{x}_1)} \mathbb{E}_{\rho(\mathbf{y}_0/\mathbf{y}_1)} \|\mathbf{x}_0 - \mathbf{y}_0\| \leq K^{T-1} \|\mathbf{x}_{T-1} - \mathbf{y}_{T-1}\| + \sum_{t=2}^{T-1} \sum_{i=1}^{t-1} K^i \mathbb{E}_{\epsilon, \epsilon} [\|\epsilon - \epsilon\|]. \quad (14)$$

Using the induction hypothesis, the linearity of the expectation, Lemma 3.3 with  $t := T$ ,

$$\begin{aligned} & \mathbb{E}_{\rho(\mathbf{x}_{T-1}/\mathbf{x}_T)} \mathbb{E}_{\rho(\mathbf{y}_{T-1}/\mathbf{y}_T)} \dots \mathbb{E}_{\rho(\mathbf{x}_0/\mathbf{x}_1)} \mathbb{E}_{\rho(\mathbf{y}_0/\mathbf{y}_1)} \|\mathbf{x}_0 - \mathbf{y}_0\| \\ & \leq \mathbb{E}_{\rho(\mathbf{x}_{T-1}/\mathbf{x}_T)} \mathbb{E}_{\rho(\mathbf{y}_{T-1}/\mathbf{y}_T)} \left[ \sum_{t=1}^{T-1} K^t \|\mathbf{x}_{T-1} - \mathbf{y}_{T-1}\| + \sum_{t=2}^{T-1} \sum_{i=1}^{t-1} K^i \mathbb{E}_{\epsilon, \epsilon} [\|\epsilon - \epsilon\|] \right] \\ & = \sum_{t=1}^{T-1} K^t \mathbb{E}_{\rho(\mathbf{x}_{T-1}/\mathbf{x}_T)} \mathbb{E}_{\rho(\mathbf{y}_{T-1}/\mathbf{y}_T)} [\|\mathbf{x}_{T-1} - \mathbf{y}_{T-1}\|] + \sum_{t=2}^{T-1} \sum_{i=1}^{t-1} K^i \mathbb{E}_{\epsilon, \epsilon} [\|\epsilon - \epsilon\|] \\ & = \sum_{t=1}^{T-1} K^t \|\mathbf{x}_T - \mathbf{y}_T\| + \sum_{t=2}^{T-1} \sum_{i=1}^{t-1} K^i \mathbb{E}_{\epsilon, \epsilon} [\|\epsilon - \epsilon\|] \\ & = \sum_{t=1}^{T-1} K^t \|\mathbf{x}_T - \mathbf{y}_T\| + \sum_{i=1}^{T-1} \sum_{t=2}^{T-1} K^i \mathbb{E}_{\epsilon, \epsilon} [\|\epsilon - \epsilon\|] \\ & = \sum_{t=1}^{T-1} K^t \|\mathbf{x}_T - \mathbf{y}_T\| + \sum_{t=2}^{T-1} \sum_{i=1}^{t-1} K^i \mathbb{E}_{\epsilon, \epsilon} [\|\epsilon - \epsilon\|]. \end{aligned}$$

□

## B Numerical Experiments

The goal of these experiments is to assess the numerical value of the bound of Theorem 3.1 on a synthetic dataset. The data-generating distribution is chosen to be the uniform distribution on the square of side 2, centered at the origin. Figure 2 shows samples from this target distribution.

The backward process uses a shared network with fully connected layers, and 128 hidden units each. The model is trained on 50,000 samples from the original distribution, and the bound is computed with  $n = 5,000$  independent samples. Samples from the trained model are shown in Figure 3.

We computed the bound for different values of  $n$ . Given that the datapoints are confined to a square of side 2, using the primal form of the Wasserstein distance yields a straightforward upper bound of  $W_1(\mu, \nu) \leq \frac{1}{\sqrt{8}} \approx 0.3535$ . We estimated the Lipschitz norms  $K^t$  using  $K_t$  from Remark 3.2, and the expected norms in the last two terms of Theorem 3.1 are estimated using  $10^6$  independent samples from each distribution.

	$n/10$	$n/5$	$n/2$	$n$	$n/0.5$	$n/0.1$
Bound value	1.124	1.231	1.5181	2.035	3.056	11.061




Figure 2: The points represent 2000 samples from the target data-generating distribution.




Figure 3: The points represent 2000 samples from the trained diffusion model.