Score-based generative modeling secretly minimizes the Wasserstein distance

Anonymous Author(s)
Affiliation
Address
email

Abstract

Score-based generative models are shown to achieve remarkable empirical performances in various applications such as image generation and audio synthesis. However, a theoretical understanding of score-based diffusion models is still incomplete. Recently, Song et al. showed that the training objective of score-based generative models is equivalent to minimizing the Kullback-Leibler divergence of the generated distribution from the data distribution. In this work, we show that score-based models also minimize the Wasserstein distance between them. Specifically, we prove that the Wasserstein distance is upper bounded by the square root of the objective function up to multiplicative constants and a fixed constant offset. Our proof is based on a novel application of the theory of optimal transport, which can be of independent interest to the society. Our numerical experiments support our findings. By analyzing our upper bounds, we provide a few techniques to obtain tighter upper bounds.

1 Introduction

The idea of gradually perturbing the data distribution to define the forward process and learning the reverse process has recently demonstrated remarkable performances in generative modeling. Score-based generative models [33, 35, 36] and diffusion probabilistic models [29, 17] have been used in numerous applications, including image generation [25, 12], and audio synthesis [9, 21].

Both model families have the following two processes. The given data distribution is incrementally perturbed by adding small noise in the forward process. This forward process can be understood as the standard diffusion process [37]. The reverse process requires the so-called *score function* [2], the gradient of the log probability density function. To approximate the score function, it is common to minimize a weighted combination of score matching losses. Once the score function is approximated, one can generate samples from randomly generated noise by applying the learned reverse process.

Note that score-based models do not explicitly minimize any sort of similarity between *the model distribution* (the distribution of the generated samples) and the data distribution. Therefore, it is unclear whether or not scored-based generative models can minimize divergence (or distance) between the model distribution and the data distribution. In very recent work [31], the authors show that optimizing the score matching loss is indeed equivalent to minimizing the upper bound on the Kullback–Leibler divergence (KL divergence). Therefore, it is expected that the model distribution is similar (in the sense of KL divergence) to the data distribution if the score matching loss is small enough.

A natural question follows – *do score-based models minimize any other divergence or distance than the KL divergence?* At first glance, this seems unlikely. For instance, Generative Adversarial Networks (GANs) [16] were shown to minimize different divergences depending on the choice of the loss function. The original GAN loss function turns out to capture the Jensen–Shannon

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

divergence [16], while the W-GAN loss function captures the Wasserstein distance [3]. We refer the readers to [10] for more details.

On the other hand, a unique aspect of score-based models is that the score function is defined in the same space where data lies in. In stark contrast, all the other popular generative models such as GANs [16], VAEs [19], and Normalizing Flows [20] use loss functions that only care about likelihoods and do not directly operate in the knowledge of the underlying data space.

In this work, we answer this question by revealing the following property of score-based modeling:

Score-based modeling minimizes the Wasserstein distance.

We note that this is quite surprising as there is no direct relationship between the KL divergence and the Wasserstein distance, in general. Therefore, this finding is not implied by the aforementioned finding.

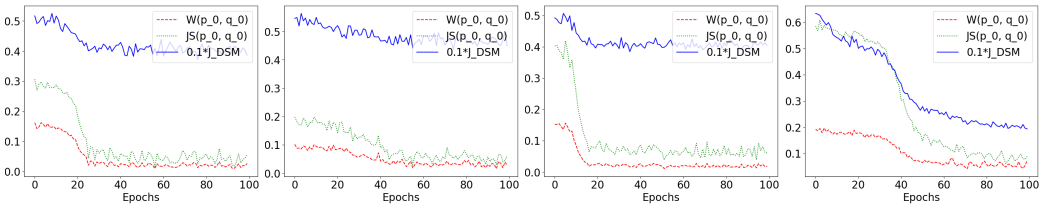

(a) `1 cluster in 1D` (b) `2 clusters in 1D` (c) `1 cluster in 2D` (d) `4 clusters in 2D`

Figure 1: Training curves of loss function J_{DSM} (eq. (14)), Wasserstein distance and JS divergence during training. The training data is sampled from mixture of 1D and 2D Gaussian clusters. (a): samples from one cluster $\mathcal{N}(0, 0.1)$. (b): samples from two clusters $\mathcal{N}(0, 0.1)$ and $\mathcal{N}(0.5, 0.1)$ with equal weights. (c): samples from one cluster $\mathcal{N}(\mathbf{0}, 0.01\mathbf{I})$. (d): samples from four clusters $\mathcal{N}((\pm 0.5, \pm 0.5)^\top, 0.01\mathbf{I})$ with equal weights.

Using the theory of optimal transport, we investigate the evolution of data distribution based on Partial differential equations (PDEs). This allows us to estimate the Wasserstein distance between the given data distribution and the distribution of the generated samples. Our work theoretically guarantees that the score-based generative models approximate data distributions in the space of probability measures equipped with the Wasserstein distance. Our main contributions and findings are summarized as follows:

- (Theorem 1 & Corollary 1) Using the theory of optimal transport, we prove that the Wasserstein distance between the data distribution and the model distribution is bounded from above. The upper bound is given as the square root of the score-matching loss J_{SM} up to multiplicative constants and a fixed constant offset. The main idea is to estimate the time-derivative of the Wasserstein metric by using the form of the continuity equations.
- (Theorem 2) We prove $J_{SM} \leq J_{DSM}$ under suitable assumptions, which gives us a more practical upper bound since J_{DSM} is much easier to estimate in practice and is widely used as a loss function to train score-based models.
- (Theorem 3) We analyze how the model distribution changes if the data distribution is perturbed initially. This guarantees that the score based-model is robust under noise.
- (Section 4) Our numerical experiments support the upper bound of our analysis. We also explore the effect on the weight decay coefficient in training and different choices of the number of total timesteps T.

Notations: $x = (x_1, x_2, \cdots, x_d)^\top \in \mathbb{R}^d$ and $t \in [0, T]$ denote the spatial variable and the time variable, respectively. Here, $\mathbb{R}^d$ is the d-dimensional Euclidean space and $T > 0$ is the time horizon. The gradient of a real-valued function ρ with respect to the spatial variable and the time-derivative of ρ are denoted by $\nabla \rho = \left(\frac{\partial \rho}{\partial x_1}, \frac{\partial \rho}{\partial x_2}, \cdots, \frac{\partial \rho}{\partial x_d}\right)^\top$, and $\partial_t \rho$, respectively. The Laplacian of ρ is given as $\Delta \rho = \nabla \cdot (\nabla \rho)$. Here, $\nabla \cdot F = \frac{\partial F_1}{\partial x_1} + \frac{\partial F_2}{\partial x_2} + \cdots + \frac{\partial F_d}{\partial x_d}$ denotes the divergence of $F = (F_1, F_2, \cdots, F_d)$ with respect to the spatial variable x. $\mathcal{N}(\mu, \Sigma)$ indicates Gaussian distribution

with mean μ and covariance Σ. $\|\cdot\|$ indicates L_2 norm. We denote the data distribution and the distribution of the generated samples by p_0 and q_0.

2 Preliminaries and related works

In this section, we analyze the forward process and the reverse process in the score-based model through Partial differential equations (PDEs). In particular, we investigate the time evolution of marginal distributions through the two processes.

2.1 The forward process

As in [37], the forward process can be modeled by the standard diffusion process $\{x(t)\}_{t=0}^{T}$ for $t \in [0, T]$ such that $x(0) \sim p_0$ and

$$dx = f(x, t)dt + g(t)d\mathbf{w}. \tag{1}$$

Here, $\mathbf{w}$ is the standard Wiener process, $f : \mathbb{R}^d \times [0, T] \to \mathbb{R}^d$ is the drift coefficient of $x(t)$, and $g : [0, T] \to \mathbb{R}$ is the diffusion coefficient of $x(t)$.

Let $p_t = p(\cdot, t)$ be a probability density function of the random variable $x(t)$ on $\mathbb{R}^d$ at time $t \in [0, T]$. Then, the function $p : \mathbb{R}^d \times [0, T] \to [0, +\infty)$ satisfies the following partial differential equations (PDEs), namely the Fokker–Planck equation or the forward Kolmogorov equation as in [4, Theorem 2.2]:

$$\partial_t p + \nabla \cdot (pf) - \frac{g^2}{2}\Delta p = 0, \;\; p(\cdot, 0) = p_0. \tag{2}$$

where the pdf p_0 represents the data distribution. Under suitable assumptions on f and g, a solution to the above initial value problem uniquely exists. We refer the readers to [5] for more details.

2.2 Understanding the reverse process through PDEs

To generate samples from $t = T$ to 0, the score-based model uses the reverse-time SDE [2]:

$$dx = (f(x, t) - g(t)^2 \nabla \log p)dt + g(t)d\bar{\mathbf{w}} \tag{3}$$

Here, $\bar{\mathbf{w}}$ is a standard Wiener process when time flows backward from $t = T$ to $t = 0$.

As in the previous section, for given $x(T) \sim p_T$, a probability density function of the random variable $x(t)$ from in the above process (3) satisfies the final value problem:

$$\partial_t p + \nabla \cdot \left(p\left(f - g^2 \nabla \log p\right)\right) + \frac{g^2}{2}\Delta p = 0, \;\; p(\cdot, T) = p_T. \tag{4}$$

Using this simple identity, $\Delta p = \nabla \cdot (\nabla p) = \nabla \cdot \left(p \frac{\nabla p}{p}\right) = \nabla \cdot (p \nabla (\log p))$, it can be seen that the differential equations in (2) and (4) are indeed identical.

2.3 The objective function

The score of the probability distribution, $\nabla \log p_t$, is approximated by a function $s_\theta : \mathbb{R}^d \times [0, T] \to \mathbb{R}^d$. For instance, in [31], the authors consider the weighted mean squared error (MSE):

$$J_{SM}(\theta; \lambda) := \frac{1}{2}\int_0^T \lambda(t)\mathbb{E}_{p_t}\left[\|\nabla \log p_t(x) - s_\theta(x, t)\|^2\right] dt \tag{5}$$

where $\lambda : [0, T] \to (0, \infty)$ is a positive weighting function.

Replacing $\nabla \log p$ in (4) with s_θ, we deduce a new reverse process. For given noise distribution q_T, a probability density function $q_t = q(\cdot, t)$ at time $t \in [0, T]$ solves

$$\partial_t q + \nabla \cdot \left(q\left(f - g^2 s_\theta\right)\right) + \frac{g^2}{2}\Delta q = 0, \;\; q(T, \cdot) = q_T. \tag{6}$$

Here, we introduced a new function $q : \mathbb{R}^d \times [0, T] \to [0, +\infty)$ to illustrate the reverse process.

2.4 Related works

Optimal transport The Wasserstein distance has a solid connection to the Fokker–Planck equation. In particular, the Fokker–Planck equation can be understood as the gradient flow in the space of probability measures equipped with the 2-Wasserstein distance [1, 18, 26]. For two probability measures μ and ν on $\mathbb{R}^d$, we define the 2-Wasserstein or Monge-Kantorovich distance as

$$W_2(\mu,\nu) := \inf\left\{\int_{\mathbb{R}^d\times\mathbb{R}^d} \|x-y\|^2 d\gamma : \gamma \in \Pi(\mu,\nu)\right\}^{\frac{1}{2}}. \tag{7}$$

Here, $\Pi(\mu,\nu)$ is the set of all couplings π of μ and ν where a Borel probability measure π on $\mathbb{R}^d\times\mathbb{R}^d$ satisfies $\pi(A\times\mathbb{R}^d)=\mu(A)$ and $\pi(\mathbb{R}^d\times B)=\nu(B)$ for all measurable subsets $A, B\subset\mathbb{R}^d$.

Under suitable assumptions on the drift coefficient and the diffusion coefficient, a solution to the Fokker–Planck equation satisfies the so-called Wasserstein contraction property [6, 7, 28], which plays an important role in our analysis. More discussions can be found in Section 3.3 and Section 5.

Score matching and diffusion probabilistic models Score matching models [34, 36] use the score function to reverse the noise corruption in a continuous time range by solving stochastic differential equations (SDEs). Diffusion probabilistic models [29, 17] adopt a similar idea of modeling tractable denoising process to reverse each discrete step of noise corruption with probabilistic models. In fact, [36] show that diffusion models are score matching models with certain forms of SDE, which unifies the whole framework. Diffusion models are gaining increasing attention: recent advances [30, 25, 12] show that diffusion models can generate even better images than GANs while enjoying a more stable training process, which encourages it to be widely applied in various tasks, such as text-to-image and image editing [24, 23], audio generation [9, 21], etc.

3 Approximation in the Wasserstein space

In this section, we estimate the Wasserstein distance between a given data distribution p_0 and the marginal of the generated samples q_0 obtained by the score-based model.

3.1 The main results

We present our main results under the following assumptions.

(A1) The drift coefficient $f:\mathbb{R}^d\times[0,T]\to\mathbb{R}^d$ is Lipschitz continuous in the space variable x: there exists a positive constant $L_f(t)\in(0,\infty)$, depending on $t\in[0,T]$, such that for all $x,y\in\mathbb{R}^d$

$$\|f(x,t)-f(y,t)\| \le L_f(t)\|x-y\|. \tag{8}$$

(A2) $s_\theta:\mathbb{R}^d\times[0,T]\to\mathbb{R}^d$ satisfies *the one-sided Lipschitz condition* [13, Definition 2.1]: there exists a constant $L_s(t)\in\mathbb{R}$, depending on $t\in[0,T]$, satisfying

$$(s_\theta(x,t)-s_\theta(y,t))\cdot(x-y) \le L_s(t)\|x-y\|^2. \tag{9}$$

The usual Lipschitz constant of s_θ satisfies the above inequality (9). Unlike the usual Lipschitz constant, *the one-sided Lipschitz constant* $L_s(t)$ is not necessarily nonnegative (see Appendix B). More technical assumptions can be found in Appendix A.

We now present a general result for any choice of loss functions. Let

$$J_I(\theta) := \int_0^T g^2(t)I(t)\mathbb{E}_{p_t}\left[\|\nabla\log p_t(x)-s_\theta(x,t)\|^2\right]^{\frac{1}{2}} dt \tag{10}$$

and an integrating factor $I:[0,T]\to(0,+\infty)$,

$$I(t) := \exp\left(\int_0^t L_f(r)+L_s(r)g^2(r)dr\right). \tag{11}$$

Theorem 1. *Let p_0 and q_0 be the data distribution and a pdf obtained by the score-based model as in Section 2. Then the Wasserstein distance between p_0 and q_0 is bounded as follows:*

$$W_2(p_0,q_0) \le J_I + I(T)W_2(p_T,q_T). \tag{12}$$

The main idea of the proofs is to estimate the derivatives of the Wasserstein distance along with the time evolution, which is illustrated in Section 5. Applying the Cauchy-Schwarz inequality, we obtain the upper bound in terms of the weighted MSE J_{SM}.

Corollary 1. *Let p_0 and q_0 be given in Theorem 1. If $\lambda(t) = g(t)^2$ for all $t \in [0, T]$, then the following inequality holds:*

$$W_2(p_0, q_0) \leq \sqrt{2\left(\int_0^T g^2(t)I(t)^2 dt\right) J_{SM}} + I(T)W_2(p_T, q_T). \tag{13}$$

We defer the proofs of Theorem 1 and Corollary 1 to Appendix A.

Remark 1. *From the above result, the Wasserstein distance converges to zero as the loss function J_{SM} goes to zero if $p_T = q_T$. However, due to the competition between $I(T)$ and $W_2(p_T, q_T)$, a general convergence result remains open (see Section 6 and Appendix G).*

Remark 2. *If $g^4 I^2 \lambda^{-1}$ is integrable in $[0, T]$, the similar inequality holds for other loss functions.*

Remark 3. *Note that the above upper bound depends on the time horizon T. Numerical experiments for different choices of T are provided in Appendix F.*

3.2 The relation between J_{SM} and J_{DSM}

The above results give the guarantee that J_{SM} can upper-bound Wasserstein distance. However, J_{SM} is generally intractable due to the unknown quantity $\nabla_x \log p_t(x)$ [32], and is rarely used in practice. Alternatively, there is a conditional version of the score-matching loss J_{DSM}, which is equivalent to J_{SM} up to a constant as shown in [38]. J_{DSM} is also widely used as a loss function to train score-based models:

$$J_{DSM}(\theta, \lambda) := \frac{1}{2}\int_0^T \lambda(t)\, \mathbb{E}_{p_0(x(0))}\, \mathbb{E}_{p_{0t}(x|x(0))}[\|s_\theta(x,t) - \nabla_x \log p_{0t}(x|x(0))\|^2] dt, \tag{14}$$

where p_{0t} indicates the conditional probability of x at step t given $x(0)$. Now we present our new result, which shows that J_{DSM} is also an upper bound of J_{SM} under certain assumption.

Theorem 2. *If p_{0t} satisfies*

$$\mathrm{Var}[\mathbb{E}[(\nabla_x \log p_{0t}(x|x(0)))^\top | x(0)]] = 0, \tag{15}$$

then we have

$$J_{SM} \leq J_{DSM} \quad \textit{and} \quad W_2(p_0, q_0) \leq \sqrt{2\left(\int_0^T g^2(t)I(t)^2 dt\right) J_{DSM}} + I(T)W_2(p_T, q_T). \tag{16}$$

Corollary 2. *Let $p(x(t)|x(t-1)) = \mathcal{N}(\sqrt{1-\beta_t}x(t-1), \beta_t I)$ as in Denoising Diffusion Probabilistic Models (DDPM) [17], which corresponds to $f(x,t) = -\frac{1}{2}\beta_t x$, and $g(t) = \sqrt{\beta_t}$. Here, $\beta_t > 0$ is the noise schedule for all t, and let $\lambda(t) = g^2(t)$. Then eq. (16) holds.*

The proofs of Theorem 2 and Corollary 2 are in Appendix D. We also provide numerical results which show $J_{SM} \leq J_{DSM}$ for DDPM in Appendix E.

Remark 4. *Using Theorem 2 and Corollary 2, we can present an upper bound of the Wasserstein distance in terms of J_{DSM}, where both $W_2(p_0, q_0)$ and J_{DSM} are tractable. This enables us to empirically verify this upper bound via DDPM in Section 4.*

3.3 Perturbation and sensitivity analysis

Let $\tilde{p}_0$ be a small perturbation of the original data distribution p_0. For a given corrupted distribution $\tilde{p}_0$, denote the probability density functions in the forward process and the reverse process by $\tilde{p}_t$ and $\tilde{q}_t$, respectively. In this section, we investigate how this perturbation changes the score-based model. In particular, we are interested in the upper bound on $W_2(q_0, \tilde{q}_0)$ with respect to $W_2(p_0, \tilde{p}_0)$.

It is well known that the Wasserstein distance between two probability distributions in the forward process can be estimated as follows: for all $t \in [0, T]$

$$W_2(p_t, \tilde{p}_t) \leq \exp\left(\int_0^t L_f(r)dr\right) W_2(p_0, \tilde{p}_0) \tag{17}$$

This is the so-called Wasserstein contraction property [8]. In the reverse process, the score function in the drift coefficients may vary and thus the above inequality does not directly give us the upper bound on $W_2(q_0, \tilde{q}_0)$. Our analysis in the previous sections allows us to estimate $W_2(q_0, \tilde{q}_0)$.

It is worth noting that the Wasserstein distance satisfies the triangle inequality, unlike other statistical measures, including the Kullback–Leibler (KL) divergence. Combining our main results with $W_2(q_0, \tilde{q}_0) \leq W_2(p_0, \tilde{p}_0) + W_2(p_0, q_0) + W_2(\tilde{p}_0, \tilde{q}_0)$, we conclude the following theorem.

Theorem 3. *Let $\tilde{p}_0$ be the corrupted data distribution and $\tilde{q}_0$ be a probability density function obtained by the score-based model starting from $\tilde{p}_0$. It holds that*

$$\begin{aligned} W_2(q_0, \tilde{q}_0) \leq W_2(p_0, \tilde{p}_0) &+ \sqrt{2\left(\int_0^T g^2(t)I(t)^2 dt\right) J_{SM}} + \sqrt{2\left(\int_0^T g^2(t)\tilde{I}(t)^2 dt\right) \tilde{J}_{SM}} \\ &+ \tilde{I}(T)W_2(p_T, q_T) + \tilde{I}(T)W_2(\tilde{p}_T, \tilde{q}_T). \end{aligned} \tag{18}$$

Here, $\tilde{J}_{SM}$ and $\tilde{I}$ denote the weighted MSE and the integrating factor from the perturbed model.

Remark 5. *The Wasserstein distance between the original data distribution p_0 and the model distribution $\tilde{q}_0$ starting from the corrupted data $\tilde{p}_0$ can also be estimated similarly.*

4 Experiments

Here we present experiments on simulation datasets to verify the upper bound in eq. (16).

4.1 Experiment settings

Datasets Here we adopt three 2D datasets for simulation: One cluster Gaussian $\mathcal{N}(\mathbf{0}, 0.1\mathbf{I})$, two moons in [27], and four clusters Gaussian mixture $\mathcal{N}((\pm 0.5, \pm 0.5)^\top, 0.01\mathbf{I})$ with equal weights for each cluster. Visualization for these datasets are in Fig 2.

For forward dynamics, we use discrete timesteps with $T = 10$. For $t \in [1, 10]$, $p(x(t)|x(t-1)) = \mathcal{N}(\sqrt{1-\beta_t}x(t-1), \beta_t I)$ as in [17], i.e., $f(x, t) = -\frac{1}{2}\beta_t x$, and $g(t) = \sqrt{\beta_t}$. For β_t, we adopt a sigmoid schedule from $\beta_1 = 10^{-5}$ to $\beta_T = 10^{-2}$.

Training and evaluation We use a 3-layer neural network as the score matching model, with ReLu nonlinearity and skip-connection at the final output. Each layer is composed of a linear layer with 128 hidden neurons, and an embedding layer for 10 timesteps. For optimizer, we use AdamW [22] with lr $= 0.001$ and weight decay coefficient 0.01. To evaluate Wasserstein distance, we use POT [15] to compute $W_2(p_0, q_0)$ at the end of each training epoch. To avoid the term of $W_2(p_T, q_T)$ in eq. (16) for verification, we first perform forward pass on sampled $x(0)$ to get $x(T)$, and perform data generation on the same $x(T)$ to approximate $q_T \approx p_T$ and $W_2(p_T, q_T) \approx 0$. For loss function, we use J_{DSM} with $\lambda(t) = g^2(t)$ and batch size = 512. To get both lower and higher losses to verify our bound, we first maximize the loss for 15 epochs, and then minimize the loss till convergence.

4.2 Results

To verify the upper bound in eq. (16), we take log on each side of the equation with $W_2(p_T, q_T) = 0$ and get:

$$\log W_2(p_0, q_0) \leq \underbrace{\frac{1}{2}\log J_{DSM}}_{\text{slope}} + \underbrace{\frac{1}{2}\log\left(\int_0^T g^2(t)I(t)^2 dt\right)}_{\text{intercept}} \tag{19}$$

We present a log-log plot to verify the slope and the intercept in the upper bound, respectively: see Fig 2. The log-log plots are obtained by evaluating Wasserstein distance and the loss function during training. The red dots are the sampled data from each dataset, the green dots are generated data after loss maximization, and the purple dots are the generated data sampled at convergence. From the log-log plots, we can observe that the theoretical line (the black line) is a valid upper bound of the empirical one (the blue dots).

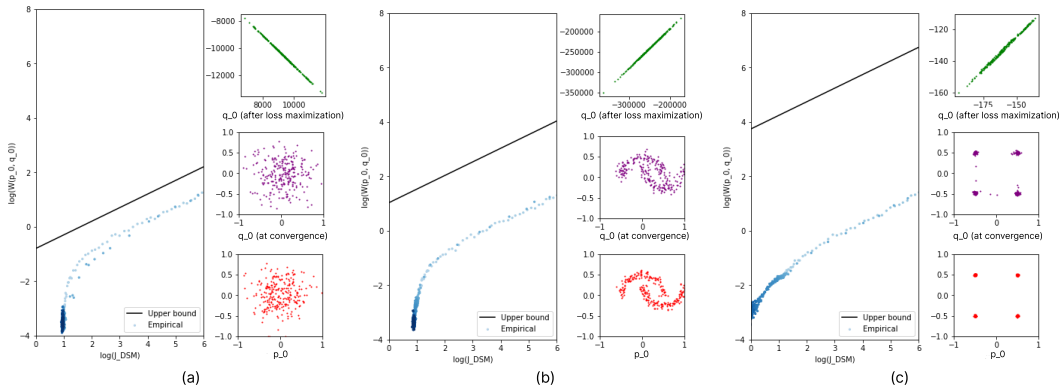

Figure 2: Log-log plots of the loss function and Wasserstein distance, with corresponding generated data at the start of the loss minimization and at convergence for each dataset. In the log-log plots, the darkness of the datapoints corresponds to training epochs: the darkest ones are at convergence.

4.3 Exploring the upper bound

Here we explore different factors that can affect the tightness of the theoretical upper bound. As shown in Fig 2, the gap between the theoretical bound and the empirical one mainly results from the intercept in eq. (19). Recall our intercept is $\frac{1}{2}\log\left(\int_0^T g^2(t)I(t)^2dt\right)$ for $I(t)$ given in (11).

4.3.1 Effect of weight regularization

Since $g(t)$ and L_f are fixed, only the one-sided Lipchitz constant $L_s(r)$ of the model s_θ is subject to the training process, which is estimated via grid search in our experiments.

Notice that the loss function J_{DSM} aims to minimize the distance between $s_\theta(x(t),t)$ and $\nabla_x \log p_{0t}(x|x(0))$. Thus, L_s at convergence is determined by dynamics $\nabla_x \log p_{0t}(x|x(0))$. If the one-sided Lipschitz constant for $\nabla_x \log p_{0t}(x|x(0))$ is small, then the intercept would be expected to be small (see example in Fig 2a of Gaussian cluster with moderately high variance). If the one-sided Lipschitz constant for $\nabla_x \log p_{0t}(x|x(0))$ is large, to minimize the loss, the intercept would also be large (see example in Fig 2c of Gaussian clusters with lower variance). However, it can be possible to make the upper bound tighter while sacrificing the training loss by constraining L_s.

To constrain L_s, the simplest and most efficient way is to constrain the norm of the weights in the score-based model. Here we present the results of training with L_2 regularization using different coefficients on the four clusters dataset, which is suffered from a large gap between empirical and theoretical bounds on the intercept. As we can observe in Fig 3a, the loss at convergence will increase when we impose stronger regularization on the weights, while the intercept decreases. Importantly, the intercept approaches the empirical one without not exceeding it, even with a very large weight decay coefficient.

We also present the log-log training plots for weight decay $= 0.01, 0.1, 0.5, 1, 5$ respectively: In Fig 4, we can observe that *as we impose stronger weight decay regularization, the theoretical upper bound will approach the empirical one, though the loss at convergence will increase*.

4.3.2 Effect of the number of total timesteps

Notice that in the definition of J_{DSM} and the intercept in eq. (19), they are both dependent on T, which might also affect the tightness of the upper bound: as T increases, the minimal $W_2(p_0, q_0)$ is

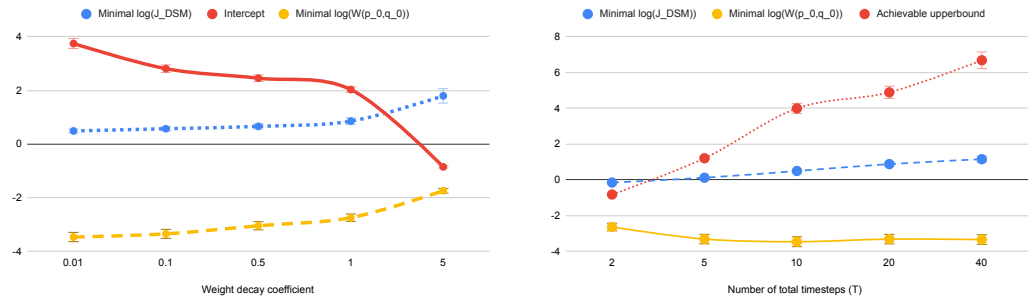

(a) Minimal losses $\log(J_{DSM})$ during training, estimation of the intercept in eq. (19) and the minimal log Wasserstein distances during training with different weight decay coefficients.

(b) Minimal losses $\log(J_{DSM})$ during training, minimal log Wasserstein distances during training and achievable upper bounds (Section 4.3.2) of different choices of T.

Figure 3: Effect of weight decay coefficient and T.

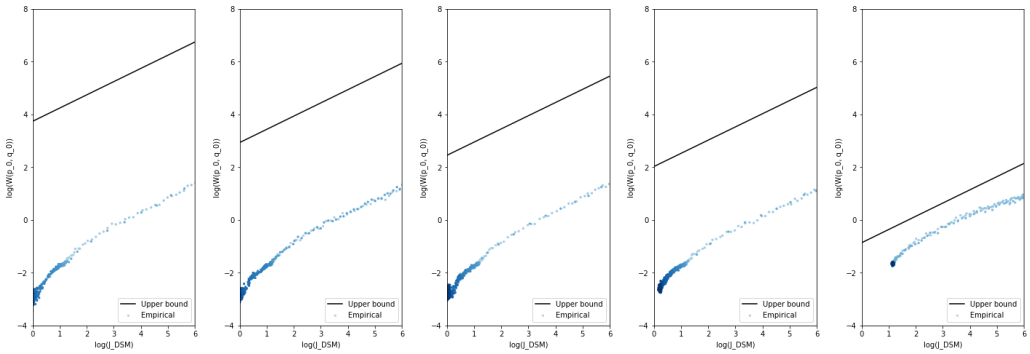

(a) Coefficient=0.01 (b) Coefficient=0.1 (c) Coefficient=0.5 (d) Coefficient=1 (e) Coefficient=5

Figure 4: Log-log plots for different weight decay coefficients. As the weight decay coefficient increases, the theoretical upper bound is approaching the empirical one.

expected to decrease, but minimal J_{DSM} would not necessarily decrease due to its dependence on T. To verify this, we conduct experiments in the four clusters dataset with $T = 2, 5, 10, 20$ and 40.

Here we first define the term *achievable upper bound* as RHS of eq. (19) when plugging in the minimal J_{DSM} in training. *As shown in Fig 3b, we can observe that when T increases, minimal $W_2(p_0, q_0)$ during training would decrease, and minimal J_{DSM} during training would increase. The achievable upper bound would also increase.* We include log-log plots of different choices of T in Appendix F, which shows that the intercept would also increase, and thus the upper bound would be looser when T increases.

5 The main idea of the proofs

We investigate the time evolution of probability measures arising in the score-based models. Using the optimal transport theory, we differentiate the Wasserstein distances, $W_2(p_t, q_t)$ with respect to time. The main idea is to express these equations (2), and (6) in the form of the continuity equation:

$$\partial_t \rho + \nabla \cdot (\rho v) = 0 \tag{20}$$

where v is a so-called velocity field. Let us denote the velocity fields of p, and q by

$$v[p] = f - \frac{1}{2} g^2 \nabla \log p, \text{ and } v[q] = \big(f - g^2 s_\theta\big) + \frac{1}{2} g^2 \nabla \log q.$$

Then, the equations (2), and (6) can be written in the form of (20) with $v[p]$ and $v[q]$, respectively.

Using these vector fields, we compute the time derivative of the Wasserstein distance between two evolving probability measures. For velocity fields $v[\rho^{(i)}]$ depending on $\rho^{(i)}$, let $\rho^{(i)}$ for $i = 1, 2$ be two solutions of the continuity equations $\partial_t \rho^{(i)} + \nabla \cdot (\rho^{(i)} v[\rho^{(i)}]) = 0$. Then, the relation between the Wasserstein distance and velocity fields arises from the theory of optimal transport:

$$\frac{1}{2}\frac{d}{dt} W_2^2(\rho_t^{(1)}, \rho_t^{(2)}) = \mathbb{E}_{\pi_t}\left[(x-y)\cdot(v[\rho_t^{(1)}](x) - v[\rho_t^{(2)}](y))\right] \tag{21}$$

where π_t is an optimal transport plan between $\rho_t^{(1)}$ and $\rho_t^{(2)}$. See [1, Theorem 8.4.7] and [28, Corollary 5.25]. This equality plays an important role when we estimate the Wasserstein distance. The following proposition is a consequence of (21) and Lemma 1 in Appendix C.

Proposition 1. *Let p_t and q_t be solutions to* (2) *and* (6)*, respectively, in* $[0, T]$*. Then,*

$$-\frac{d}{dt} W_2(p_t, q_t) \leq (L_f + L_s g^2) W_2(p_t, q_t) + g^2 b^{\frac{1}{2}} \text{ where } b(t) := \mathbb{E}_{p_t}\left[\|\nabla \log p_t(x) - s_\theta(x,t)\|^2\right]. \tag{22}$$

6 Discussions

Convergence of the Wasserstein distance If $p_T = q_T$, then our result yields that the Wasserstein distance $W_2(p_0, q_0)$ converges to zero as J_{SM} goes to zero. However, if p_T is not equal to q_T, then we find that it is not easy to obtain such a convergence result. This is due to the nature of the Langevin dynamics in the reverse process. Typically, additional assumptions such as the log-concavity are required to guarantee the Langevin dynamics convergence, for instance, [11].

In addition, the dependence of the neural network s_θ makes us hard to evaluate the precise upper bound. We expect that appropriate assumptions on s_θ and the initial distribution p_0 may resolve this issue, but we do not pursue them in this paper. See our example in Appendix G.

Estimating L_s Estimating our theoretical bound requires estimating the one-sided Lipschitz constant L_s, which is NP-hard if done via grid search [39], making it difficult to estimate in high dimensional data. Note that the two-sided Lipschitz constant is an upper bound of the one-sided one, which can be more efficiently estimated as shown in [14]. However, it would result in a looser theoretical upper bound if we adopt the two-sided Lipschitz constant, and estimating the one-sided Lipschitz constant does not enjoy the semi-definiteness of estimating the two-sided one in [14].

7 Conclusion

In this paper, we mainly discuss the relationship between the loss function of score-based models and the Wasserstein distance. With the help of the theory in optimal transport, we present a novel upper bound on Wasserstein distance in terms of the loss function used in practice. We also provide numerical experiments to verify our theoretical bound with various simulation datasets. Moreover, we explore factors that affect the estimation of the upper bound and show the corresponding influence on the training process.

References

[1] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows in metric spaces and in the space of probability measures.* Second edition. Lecture notes in Mathematics ETH Zürich. Birkhäuser Verlag, Bassel, (2008).

[2] B. D. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.

[3] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017.

[4] R. N. Bhattacharya and E. C. Waymire. *Stochastic processes with applications*. SIAM, 2009.

[5] V. I. Bogachev, N. V. Krylov, M. Röckner, and S. V. Shaposhnikov. *Fokker-Planck-Kolmogorov Equations*, volume 207. American Mathematical Soc., 2015.

[6] F. Bolley and J. A. Carrillo. Nonlinear diffusion: geodesic convexity is equivalent to wasserstein contraction. *Communications in Partial Differential Equations*, 39(10):1860–1869, 2014.

[7] J. A. Carrillo, R. J. McCann, and C. Villani. Contractions in the 2-wasserstein length space and thermalization of granular media. *Archive for Rational Mechanics and Analysis*, 179(2):217–263, 2006.

[8] J. A. Carrillo, R. J. McCann, and C. Villani. Contractions in the 2-wasserstein length space and thermalization of granular media. *Archive for Rational Mechanics and Analysis*, 179(2):217–263, 2006.

[9] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan. Wavegrad: Estimating gradients for waveform generation. *arXiv preprint arXiv:2009.00713*, 2020.

[10] I. Csiszár and P. C. Shields. Information theory and statistics: A tutorial. 2004.

[11] A. S. Dalalyan. Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(3):651–676, 2017.

[12] P. Dhariwal and A. Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34, 2021.

[13] T. Donchev and E. Farkhi. Stability and euler approximation of one-sided lipschitz differential inclusions. *SIAM journal on control and optimization*, 36(2):780–796, 1998.

[14] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas. Efficient and accurate estimation of lipschitz constants for deep neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.

[15] R. Flamary, N. Courty, A. Gramfort, M. Z. Alaya, A. Boisbunon, S. Chambon, L. Chapel, A. Corenflos, K. Fatras, N. Fournier, L. Gautheron, N. T. Gayraud, H. Janati, A. Rakotomamonjy, I. Redko, A. Rolet, A. Schutz, V. Seguy, D. J. Sutherland, R. Tavenard, A. Tong, and T. Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.

[16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[17] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.

[18] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of the Fokker–Planck equation. *SIAM journal on Mathematical Analysis*, 29(1):1–17, 1998.

[19] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[20] I. Kobyzev, S. J. Prince, and M. A. Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):3964–3979, 2020.

[21] M. W. Lam, J. Wang, D. Su, and D. Yu. Bddm: Bilateral denoising diffusion models for fast and high-quality speech synthesis. *arXiv preprint arXiv:2203.13508*, 2022.

[22] I. Loshchilov and F. Hutter. Fixing weight decay regularization in adam. 2018.

[23] C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.-Y. Zhu, and S. Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2021.

[24] A. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.

[25] A. Q. Nichol and P. Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021.

[26] F. Otto. The geometry of dissipative evolution equations: the porous medium equation. 2001.

[27] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[28] F. Santambrogio. *Optimal Transport for Applied Mathematicians*. Calculus of Variations, PDEs, and Modeling. Birkhäuser, Oct. 2015.

[29] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.

[30] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

[31] Y. Song, C. Durkan, I. Murray, and S. Ermon. Maximum likelihood training of score-based diffusion models. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 1415–1428. Curran Associates, Inc., 2021.

[32] Y. Song, C. Durkan, I. Murray, and S. Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems*, 34, 2021.

[33] Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019.

[34] Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019.

[35] Y. Song and S. Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020.

[36] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

[37] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.

[38] P. Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.

[39] A. Virmaux and K. Scaman. Lipschitz regularity of deep neural networks: analysis and efficient estimation. *Advances in Neural Information Processing Systems*, 31, 2018.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...

(a) Did you state the full set of assumptions of all theoretical results? [Yes]
(b) Did you include complete proofs of all theoretical results? [Yes]

3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]