Convergence of Consistency Model with Multistep Sampling under General Data Assumptions

Yiding Chen¹ Yiyi Zhang¹ Owen Oertell¹ Wen Sun¹

Abstract

Diffusion models accomplish remarkable success in data generation tasks across various domains. However, the iterative sampling process is computationally expensive. Consistency models are proposed to learn consistency functions to map from noise to data directly, which allows one-step fast data generation and multistep sampling to improve sample quality. In this paper, we study the convergence of consistency models when the selfconsistency property holds approximately under the training distribution. Our analysis requires only mild data assumption and applies to a family of forward processes. When the target data distribution has bounded support or has tails that decay sufficiently fast, we show that the samples generated by the consistency model are close to the target distribution in Wasserstein distance; when the target distribution satisfies some smoothness assumption, we show that with an additional perturbation step for smoothing, the generated samples are close to the target distribution in total variation distance. We provide two case studies with commonly chosen forward processes to demonstrate the benefit of multistep sampling.

1. Introduction

Diffusion models have been widely acknowledged for their high performance across various domains, such as material and drug design (Xu et al., 2022; Yang et al., 2023; Xu et al., 2023), control (Janner et al., 2022), and text-to-image generation (Black et al., 2023; Oertell et al., 2024). The key idea of diffusion models is to transform noise into approximate samples from the target data distribution by iterative denoising. This iterative sampling process typically involves numerical solutions of SDE or ODE, which is computationally expensive especially when generating high-resolution images (Ho et al., 2020; Song et al., 2021; Lu et al., 2022; Zhang & Chen, 2023; Song & Dhariwal, 2024).

Consistency model (CM) (Song et al., 2023) is proposed to accelerate sample generation by learning a consistency function that maps from noise to data directly. It allows both one-step fast data generation and multistep sampling to trade computation for sample quality. Consistency model can be trained with consistency distillation or consistency training (Song et al., 2023), which enforce that any points on the same trajectory specified by the probability-flow ODE are mapped to the same origin, i.e. the *self-consistency* property. Despite the empirical success of consistency models, their theoretical foundations remain inadequately understood. In particular, recent studies (Luo et al., 2023; Song & Dhariwal, 2024; Kim et al., 2024) observe diminishing improvements in sample quality when increasing the number of steps in multistep sampling. They find that two-step generation enhances the sample quality considerably while additional sampling steps provide minimal improvements. Such phenomenon motivates the theoretical understanding on consistency models, especially on multistep sampling.

The analysis of consistency models can be challenging for the following reasons:

Mismatch on the initial starting distributions: Consistency models generate samples from Gaussian noise (Song et al., 2023) while the ground truth reverse processes (i.e., the denoising process) start from the marginal distribution of the forward process, which is unknown in practice. As a consequence, we need to analyze the error caused by the mismatch in starting distributions. This difficulty shows up even if we have access to the ground truth consistency function: the consistency function is not Lipschitz even for distributions as simple as Bernoulli, which makes it challenging to analyze this error pointwise. Because the consistency function is the solution to the probability flow ODE, it is natural to consider the stability of the initial value problem. However, without a strong assumption on the consistency function, this approach results in an upper bound with exponential dependency in problem parameters (see Appendix A).

¹Cornell University. Correspondence to: Yiding Chen <yc2773@cornell.edu>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

Approximate self-consistency: While the training process enforces the self-consistency property, it is impractical to obtain a consistency function estimate with the point-wise exact self-consistency due to various error sources during training (e.g., optimization errors, statistical errors from finite training examples). It is thus natural to focus on the case where the consistency estimator only has approximate self-consistency under the training distribution. The key challenge is how to transfer the approximate self-consistency measured under the training distribution to the quality of the generated samples (e.g., Wasserstein distance between the learned distribution and the ground truth distribution).

Complexity of multistep sampling: We still have limited understanding of the theoretical advantages of performing multistep sampling during inference steps of CM. When performing multistep inference, we need to apply the consistency estimator repeatly to the distributions that are different from its original training distribution. Since we can only guarantee approximate self-consistency under the training distribution, analyzing the benefit of multistep sampling requires us to carefully bound the divergence between the training distributions and the test distributions where consistency estimator will be applied during inference time.

1.1. Our Contributions

In this paper, we analyze the convergence of consistency models under minimal assumptions and provide guarantees in both Wasserstein distance and Total Variation distance.

Guarantees in Wasserstein distance. Given an arbitrary distribution with *bounded support*, our main theorem establishes guarantees in *Wasserstein* distance for *multistep sampling* with a *general set of forward processes* and an *approximate self-consistent* consistency function estimator. The assumption in our result is much weaker than those in previous works. Previous works make an implicit assumption on both the target distribution and the forward process. They assume the ground truth consistency function to be Lipschitz. In contrast, our result only assumes that the target distribution has bounded support. In addition, our analysis applies to a broad class of forward processes that captures both Variance Preserving and Variance Exploding SDEs as forward processes, whereas previous work has focused only on the former.

For illustration purposes, we summarize the instantiation of our main result applied to the Ornstein-Uhlenbeck (OU) process:

Theorem 1 (informal, see Theorem 2 and Corollary 1). Suppose the consistency function estimate is ϵ^2 accurate and the support of the target distribution P_{data} is bounded by *R*, then one-step sampling returns a distribution that is $\left(\epsilon \log \frac{R^3}{\epsilon^2}\right)$ -close to P_{data} in W_2 distance; two-step sampling returns a distribution that is $\left(\epsilon \log \frac{R^2}{\epsilon}\right)$ -close to P_{data} in W_2 distance.

Our error guarantee scales peacefully in problem parameters and is dimension-free. It shows that two-step sampling reduces the error by half in the ideal case ($\epsilon \ll 1$). Our analysis also suggests that further improvements with multistep sampling is unclear even with an increased number of sampling steps. This observation is consistent with the findings from empirical studies. Additionally, the bounded support assumption can be further relaxed to a *light-tail condition*.

Guarantees in Total Variation distance. Standard consistency model cannot guarantee closeness in Total Variation (TV) distance, as the consistency loss enforces pointwise distance, which differs fundamentally from the structure of TV distance. However, we show that incorporating an additional smoothing step provides a guarantee in TV when the target distribution meets certain smoothness assumptions.

1.2. Related Work

The theory of diffusion models has been widely studied. Chen et al. (2023b), Lee et al. (2023), and Chen et al. (2023a) study the convergence of score-based generative model and provide polynomial guarantees without assuming log-concavity or a functional inequality on the data. Our data assumption is similar to that of Lee et al. (2023), which is quite minimal. Recently, deterministic samplers with probability-flow ODE have been explored from the theoretical perspective (Li et al., 2023; Chen et al., 2024; Li et al., 2024a).

Consistency model, which learns a direct mapping from noise to data via trajectory of probability-flow ODE, is proposed to accelerate the sampling step (Song et al., 2023). Song et al. (2023) provides asymptotic theoretical results on consistency models. At a high level, they show that if the consistency distillation objective is minimized, then the consistency function estimate is close to the ground truth. However, they assume the consistency function estimator achieves exact self-consistency in a point-wise manner. Such a point-wise accurate assumption is not realistic and cannot even be achieved in a standard supervised learning setting.

Lyu et al. (2023), Li et al. (2024b), and Dou et al. (2024) provide the first set of theoretical results towards understanding consistency models. Lyu et al. (2023) shows that with small consistency loss, consistency model generates samples that are close to the target data distribution in Wasserstein distance or in total variation distance after modification. Li et al. (2024b) focuses on consistency training. Dou et al. (2024) provides the first set of statistical theory for consistency models. However, we notice that all of these works require a strong assumption on the data distribution. Specifically, they assume that the ground truth consistency function is Lipschitz. While the Lipschitz condition allows a direct approach to control the error of mismatch on the initial starting distribution, it's unclear how large the Lipschitz coefficient is. Intuitively, when the target distribution has two modes, the derivative of the ground truth consistency function can become large near the intersection of the pre-images of the two modes. This occurs because a small perturbation in the input can cause the output of the consistency function to shift from one mode to the other. This intuition is formalized in Appendix A. To overcome this, we use the data-processing inequality, which only requires approximate self-consistency and minor assumptions on target data distribution. Moreover, our upper bound is polynomial in problem parameters. Additionally, all of these works focus only on variance preserving SDEs while our results apply to a general family of forward processes.

2. Preliminaries

Score-based generative models (Song et al., 2021) and consistency models (Song et al., 2023) aim to sample from an unknown *data distribution* P_{data} in \mathbb{R}^d . We review some basic concepts and introduce relevant notations in this section.

Score-based generative model: A score-based generative model, or diffusion model (Ho et al., 2020; Song et al., 2021) defines a *forward process* $\{\mathbf{x}_t\}_{t \in [0,T]}$ by injecting Gaussian noise into the data distribution P_{data} in *d*-dimensional space \mathbb{R}^d , where $\mathbf{x}_0 \sim P_{\text{data}}$ and T > 0. In this paper, we focus on a general family of forward processes characterized by stochastic differential equations (SDEs) with the following form:

$$d\mathbf{x}_t = h(t)\mathbf{x}_t dt + g(t)d\mathbf{w}_t, \quad \mathbf{x}_0 \sim P_{\text{data}}, \qquad (1)$$

where \mathbf{w}_t is the standard Wiener process. It is known that the marginal distribution of \mathbf{x}_t in (1) is Gaussian conditioning on \mathbf{x}_0 (Kingma et al., 2021; Lu et al., 2022):

$$\mathbf{x}_t | \mathbf{x}_0 \sim \mathcal{N}(\alpha_t \mathbf{x}_0, \sigma_t^2 I), \quad \forall t \in [0, T],$$

where $\alpha_t, \sigma_t \in \mathbb{R}^+$ is specified by $h(t) = \frac{d \log \alpha_t}{dt}$, $g^2(t) = \frac{d \sigma_t^2}{dt} - 2 \frac{d \log \alpha_t}{dt} \sigma_t^2$ with proper initial conditions. α_t and σ_t^2 specify the noise schedule of the forward process. The noise schedule $\{(\alpha_t, \sigma_t^2)\}_{t \in [0,T]}$ and initial data distribution determine the marginal distribution of the forward process $\{P_t\} \in [0,T]$, where $\mathbf{x}_t \sim P_t$ and $P_0 = P_{\text{data}}$. We use $\{p_t\}_{t \in [0,T]}$ to denote the probability density functions (PDFs) of $\{P_t\}_{t \in [0,T]}$.

The forward process specified by (1) converges to Gaussian distribution $\mathcal{N}(0, \sigma_t^2 I)$ for some properly chosen $h(\cdot)$ and

 $g(\cdot)$ (Bakry et al., 2014; Song et al., 2021) (interested readers may refer to Lemma 3 for an explicit dependency on the noise schedule). The convergence of the forward process facilitates a procedure to generate samples from P_{data} , approximately: generate a sample from $\mathcal{N}(0, \sigma_T^2 I)$ and feed it to an approximate reversal of (1). However, the reverse-time SDE of (1) is usually computationally expensive.

It is known that the following *probability flow ordinary differential equation* (PF-ODE) generates the same distributions as the marginal distribution of (1) (Song et al., 2021):

$$\frac{\mathrm{d}\mathbf{x}_t}{\mathrm{d}t} = h(t)\mathbf{x}_t - \frac{1}{2}g^2(t)\nabla\log p_t(\mathbf{x}_t), \quad \mathbf{x}_0 \sim P_{\text{data}}.$$
 (2)

The time-reversal of (2) defines a deterministic mapping from noise to data, which facilitates *consistency model* (Song et al., 2023) as a computationally efficient one-step sample generation.

Consistency models: A consistency model (Song et al., 2023) is an alternative approach to generate samples from P_{data} : instead of solving the reversal of the SDE in (1), one could directly learn a *consistency function* that maps a point on a trajectory of (2) to its origin. For any **x** and $t_0 \ge 0$, let $\{\varphi(t; \mathbf{x}, t_0)\}_{t \in [0,T]}$ be the trajectory specified by (2) and *initial condition* $\mathbf{x}_{t_0} = \mathbf{x}$.¹ The (ground truth) consistency function of (2) is defined as:²

$$f^{\star}(\mathbf{x},t) := \varphi(0;\mathbf{x},t), \quad \forall \mathbf{x} \in \mathbb{R}^d, t \ge 0.$$
(3)

A consistency function enjoys the *self-consistency* property: if (\mathbf{x}, t) and (\mathbf{x}', t') are on the same trajectory of (2), they are mapped to the same origin, i.e. $f^*(\mathbf{x}, t) = f^*(\mathbf{x}', t')$.³

The self-consistency property of the ground truth consistency function $f^*(\cdot, \cdot)$ enlightens the training for consistency function via enforcing the self-consistency property instead of learning the mapping from noise to data directly. At a high level, in the training stage, we first discretize the interval [0, T] with the following partition:

$$T: 0 = \tau_0 < \tau_1 < \tau_2 < \dots < \tau_M = T.$$

¹Specifically, $\varphi(\cdot; \mathbf{x}, t_0)$ is the solution to the ODE *initial value* problem specified by (2) and $\mathbf{x}_{t_0} = \mathbf{x}$

²(Song et al., 2023) stops at time $t = \delta$ for some small $\delta > 0$ and accepts $\hat{f}(\mathbf{x}, t) = \hat{\varphi}(\delta; \mathbf{x}, t)$, an estimate for $\varphi(\delta; \mathbf{x}, t)$ as the approximate samples to avoid numerical instability. In this paper, we ignore this numerical issue to obtain a cleaner theoretical analysis.

³At a high level, this can be shown by contradiction: suppose (\mathbf{x}', t') lies on the trajectory of (\mathbf{x}, t) , meaning $\varphi(\cdot; \mathbf{x}, t)$, the trajectory of (\mathbf{x}, t) and $\varphi(\cdot; \mathbf{x}', t')$, the trajectory of (\mathbf{x}', t') intersect at (\mathbf{x}', t') . Then both trajectories satisfy the initial condition that takes value \mathbf{x}' at time t'. By Picard's existence and uniqueness theorem, the trajectories of $\varphi(\cdot; \mathbf{x}, t)$ and $\varphi(\cdot; \mathbf{x}', t')$ are identical and have the same origin.

For simplicity, we assume the partition is equal, i.e. there exists $\Delta \tau > 0$, s.t. $\tau_i = \Delta \tau \cdot i$, for i = 1, ..., M. (We defer the adaption to non-uniform distization to Appendix G) We then enforce the self-consistency property on each partition point by finding some $\hat{f}(\cdot, \cdot)$, s.t.

$$\mathbb{E}_{\mathbf{x}_{\tau_i} \sim P_{\tau_i}} \left[\left\| \hat{f}(\mathbf{x}_{\tau_i}, \tau_i) - \hat{f}(\varphi(\tau_{i+1}; \mathbf{x}_{\tau_i}, \tau_i), \tau_{i+1}) \right\|_2^2 \right]$$
(4)

is small for all i = 0, 1, ..., M - 1. This strategy is justified by our theoretical results in Section 3: even if the self-consistency property is violated slightly, the consistency function estimation will produce high-quality samples. In practice, the trajectories of the PF-ODE (2) are unknown, so the self-consistency objective cannot be optimized directly. With this regard, consistency distillation, which utilizes a pre-trained score function estimate, and consistency training, which builds an unbiased estimate for the score function, are proposed to approximate the transition on the trajectories of the PF-ODE. Interested readers can find the details in (Song et al., 2023).

Multistep sampling: Given a consistency model estimate $\hat{f}(\cdot, \cdot)$, we could generate approximate samples by feeding Gaussian noise into $\hat{f}(\cdot, \cdot)$ using *single-step* or *multistep* sampling. Given $\hat{\mathbf{x}}_T \sim \mathcal{N}(0, \sigma_T^2 I)$, one can generate sample in a single step by calculating $\hat{f}(\hat{\mathbf{x}}_T, T)$. Furthermore, one can also design a sequence of time steps by selecting $N \geq 1$ steps in the training partition \mathcal{T} :

$$T = t_1 > t_2 > \dots > t_N > 0,$$
 (5)

We refer to the sequence $\{t_i\}_{i=1:N} \subseteq \mathcal{T} \setminus \{0\}$ as sampling time schedule. Given this sampling time schedule, one can alternatingly denoise by calculating $\hat{\mathbf{x}}_0^{(i)} = \hat{f}(\hat{\mathbf{x}}_{t_i}^{(i)}, t_i)$ and inject noise by drawing $\hat{\mathbf{x}}_{t_{i+1}}^{(i+1)} \sim \mathcal{N}(\alpha_{t_{i+1}}\hat{\mathbf{x}}_0^{(i)}, \sigma_{t_{i+1}}^2 I)$, where $\hat{\mathbf{x}}_{t_1}^{(1)} = \hat{\mathbf{x}}_T \sim \mathcal{N}(0, \sigma_T^2 I)$ and $i = 1, \ldots, N$. The $\hat{\mathbf{x}}_0^{(N)}$ in the last step is the output of the sampling process. When N = 1, this degenerates to single-step sampling. For completeness, we summarize this process in Algorithm 1 in Section B. For a concise presentation, we defines $\left\{\hat{P}_{t_i}\right\}_{i=1:N}$ to be the sequence of marginal distributions of $\left\{\hat{\mathbf{x}}_{t_i}^{(i)}\right\}_{i=1:N}$ to be the sequence of marginal distributions of $\left\{\hat{\mathbf{x}}_{0}^{(i)}\right\}_{i=1:N}$. In the following, we may reuse $\hat{f}(\cdot, \cdot)$ for operation on distributions. Specifically, for any distribution P and $t \ge 0$, we use $\hat{f}(P,t)$ to denote the distribution of $\hat{f}(\mathbf{x},t)$ when $\mathbf{x} \sim P$. In Section 3, we study how multistep sampling influences the sample quality from the theoretical perspective.

Performance metric: In this paper, we study the *sample quality* generated by a consistency function estimate $\hat{f}(\cdot, \cdot)$

and the multistep sampling procedure introduced above. To quantify the sample quality, we establish upper bounds on 2-Wasserstein distance (W_2) in Euclidean norm, and upper bounds on Total Variation (TV) distance. The 2-Wasserstein distance between two distributions P and Q is defined as:

$$W_2(P,Q) := \inf_{\gamma \in \Gamma(P,Q)} \sqrt{\mathbb{E}_{(\mathbf{x},\mathbf{y}) \sim \gamma} \left[\|\mathbf{x} - \mathbf{y}\|_2^2 \right]},$$

where $\Gamma(P,Q)$ is the set of all joint distributions such that the marginal distribution over the first random variable is P and the marginal distribution over the second random variable is Q.

Total Variation distance between two distributions P and Q is defined as:

$$TV(P,Q) := \frac{1}{2} \|p(\mathbf{x}) - q(\mathbf{x})\|_{1},$$

where $p(\cdot)$ is the PDF of P and $q(\cdot)$ is the PDF of Q.

3. Main Results

In this section, we present theoretical guarantees on sample quality for consistency models with multistep sampling. We first present two sets of results for the *general* forward process in (1) with *arbitrary* sampling time schedule: in Section 3.1, we demonstrate that the generated samples are close to the target data distribution P_{data} in W_2 when P_{data} has bounded support or satisfies some tail condition; with an additional smoothing step, we show guarantee in TV distance for P_{data} with smoothness condition in Section 3.2. To illustrate the general results and gain better understanding on the multistep sampling, we choose two special SDEs as forward processes and design sampling time schedules in Section 3.3.

The natural central assumption in our theoretical results is a good consistency function estimate:

Assumption 1 (A proper consistency model). Suppose $\hat{f}(\mathbf{x}, 0) = \mathbf{x}$ for all $\mathbf{x} \in \mathbb{R}^d$ and there exists $\epsilon_{cm} > 0$, s.t. (4) $\leq \epsilon_{cm}^2$ for all i = 0, 1, ..., M - 1.

The condition related to the accuracy of the consistency function estimate is necessary: we cannot generate good samples with an arbitrary function. Instead of assuming the output of $\hat{f}(\cdot, \cdot)$ and $f^*(\cdot, \cdot)$ to be close directly, we only require the self-consistency property to hold *approximately* under its training distribution, which aligns with the objective function when training for $\hat{f}(\cdot, \cdot)$. Note that our assumption does not imply \hat{f} will be self-consistent in a point-wise manner.

The self-consistency objective (4) can be approximated via *consistency distillation* or *consistency training* (Song et al., 2023). Consistency distillation uses a pre-trained score function (an estimation for $\nabla \log p_t(\cdot)$) to approximate $\varphi(\cdot; \cdot, \cdot)$

and train for $\hat{f}(\cdot, \cdot)$ with target network and online network. In Section F, we incorporate consistency distillation with minor modifications into our framework without additional data assumptions. On the other hand, consistency training constructs an unbiased estimator for $\nabla \log p_t(\mathbf{x}_t)$ to approximate (4). Theorem 2 of Song et al. (2023) shows that the self-consistency loss (4) can be approximated by consistency training under proper conditions when $\Delta \tau$ is small.

In (4), we use $\|\cdot\|_2^2$ as an error metric, which agrees with the choice in practice (Luo et al., 2023; Song et al., 2023). The metric $\|\cdot\|_2^2$ aligns better with the theoretical analysis: on the one hand, Lemma 2 demonstrates that this metric translates naturally to the 2-Wasserstein metric W_2 ; on the other hand, $\|\cdot\|_2^2$ is more suitable for the multi-step sampling because the squared error contracts nicely in the forward process with Gaussian noise as shown by Lemma 1 and 3.

3.1. Guarantees in Wasserstein Metric

We now provide upper bounds on the sampling error in W_2 distance.

Theorem 2 (W_2 error for distributions with bounded support). Suppose Assumption 1 holds. Suppose there exists R > 0, s.t. $\sup_{\mathbf{x} \in \operatorname{supp}(P_{data})} \|\mathbf{x}\|_2 \leq R$ and $\|\hat{f}(\mathbf{x},t)\|_2 \leq R$ for all $(\mathbf{x},t) \in \mathbb{R}^d \times [0,T]$, Let $\hat{P}_0^{(N)}$ be the output of *N*-step sampling. Then $W_2(\hat{P}_0^{(N)}, P_{data})$, the error in W_2 is upper bounded by:

$$2R\left(\underbrace{\frac{\alpha_{t_1}^2}{4\sigma_{t_1}^2}R^2}_{(i)} + \underbrace{\sum_{j=2}^N \frac{\alpha_{t_j}^2}{4\sigma_{t_j}^2} t_{j-1}^2 \frac{\epsilon_{\rm cm}^2}{\Delta\tau^2}}_{(ii)}\right)^{1/4} + \underbrace{t_N \frac{\epsilon_{\rm cm}}{\Delta\tau}}_{(iii)}, \quad (6)$$

where $\epsilon_{\rm cm}$ is the consistency loss. R is the diameter of distribution. α_t, σ_t are the drift and variance factor in the forward diffusion process. $\Delta \tau$ is the time step size according to the partition.

Compared to $P_{\text{data}} = f^*(P_{t_N}, t_N)$, the sampling error of $\hat{P}_0^{(N)} = \hat{f}(\hat{P}_{t_N}, t_N)$ comes from: (i). the error of starting from Gaussian distribution $\mathcal{N}(0, \sigma_T^2 I)$ instead of P_T ; (ii) the error accumulated in the previous sampling steps; (iii). using an inaccurate consistency function estimate $\hat{f}(\cdot, \cdot)$ instead of $f^*(\cdot, \cdot)$. The term $\frac{\alpha_t^2}{\sigma_t^2}$ characterizes the convergence of the forward process as demonstrated by Lemma 3. It converges to 0 quickly for reasonable forward SDE (1). Asymptotically, the right hand side of (6) goes to 0 as $t_1 \to \infty$ and $\epsilon_{\rm cm} \to 0$.

One implication from (6) is: increasing the number of sampling steps does not necessarily improve the sample quality. When using more sampling steps: on one hand, $\frac{\alpha_{t_1}^2}{4\sigma_{t_2}^2}R^2$ + $\sum_{j=2}^{N} \frac{\alpha_{t_i}^2}{4\sigma_{t_i}^2} t_{i-1}^2 \frac{\epsilon_{cm}^2}{\Delta \tau^2}$, an upper bound on KL $(P_{t_N} \parallel \hat{P}_{t_N})$,⁴ accumulates; on the other hand, $t_N \frac{\epsilon_{cm}}{\Delta \tau}$, the error from an inaccurate consistency function decreases due to a shorter t_N . The design of sampling time schedule $\{t_i\}_{i=1:N}$, which depends on the noise schedule $\{(\alpha_t, \sigma_t^2)\}_t$, is crucial in achieving good sample quality. We defer design choices for some specific forward processes and simplified upper bounds to Section 3.3. Indeed, both our simulation in Appendix I and the results in Luo et al. (2023) provide empirical evidence that increasing the number of sampling steps can harm the sample equality.

When $\Delta \tau$ decreases, on the one hand, there would be more intermediate steps in the error decomposition of the consistency function estimate given a fix t (see Lemma 2); on the other hand, using a smaller $\Delta \tau$ allows a smaller t_N and may potentially decrease $\epsilon_{\rm cm}$ as well.

The technique in Theorem 2 can be extended to distributions without finite support but with proper tail conditions. The detailed discussion is presented in Appendix D.

3.2. Guarantee in Total Variation Distance

In the sampling process of consistency models, it is nontrivial to control the error in TV distance. This difficulty arises even when we sample with a single step and have access to the exact marginal distribution P_T . Assumption 1 ensures that $\hat{f}(P_T,T)$ is close to $f^*(P_T,T)$ in W_2 . However, W_2 and TV have very different structures: W_2 controls the pointwise distance between distributions while TV only focuses on the density of the distribution. Even if $W_2(\hat{P}_0^{(N)}, P_{data})$ is small, the densities of $\hat{f}(P_T,T)$ and $f^*(P_T,T)$ may not overlap well (see Figure 1a) and $TV(\hat{f}(P_T,T), f^*(P_T,T))$ can be as large as 1 if $\hat{f}(P_T,T)$ is nearly deterministic while $f^*(P_T,T)$ has large variance. As a result, it's not possible to control TV distance only with conditions on W_2 distance in general.

One solution is to perturb $\hat{P}_0^{(N)}$ slightly with Gaussian noise $\mathcal{N}(0, \sigma_{\epsilon}^2)$. With this perturbation, $\hat{P}_0^{(N)} * \mathcal{N}(0, \sigma_{\epsilon}^2)$ and $P_{\text{data}} * \mathcal{N}(0, \sigma_{\epsilon}^2)$ could have better overlap and be closer in TV (see Figure 1b), where we use P * Q to denote the *convolution* of distribution P and Q. When P_{data} satisfies smoothness assumption, the perturbation will not change P_{data} too much so $\text{TV}(P_{\text{data}} * \mathcal{N}(0, \sigma_{\epsilon}^2), P_{\text{data}})$ is small (See Figure 1c).

Theorem 3 (TV error for distributions under smoothness assumption). Suppose Assumption 1 holds. Let $p_{data}(\cdot)$ be the PDF of P_{data} . If $\log p_{data}(\cdot)$ is L-smooth, then for all $\sigma_{\epsilon} > 0$, the error in TV distance of the smoothed output, i.e.

⁴We use KL($P \parallel Q$) to denote the Kullback–Leibler (KL) divergence of distribution P from distribution Q, which is defined as: KL($P \parallel Q$) := $\int_{\mathbf{x} \in \mathbb{R}^d} p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} d\mathbf{x}$.



(a) TV distance between $\hat{P}_0^{(N)}$ and P_{data} .



(b) TV distance between $\hat{P}_0^{(N)} * \mathcal{N}(0, \sigma_{\epsilon}^2 I)$ (c) TV distance between $P_{\text{data}} * \mathcal{N}(0, \sigma_{\epsilon}^2 I)$ and $P_{\text{data}} * \mathcal{N}(0, \sigma_{\epsilon}^2 I)$). and P_{data} .

Figure 1: Smoothing by additional perturbation

 $\operatorname{TV}(\hat{P}_0^{(N)} * \mathcal{N}(0, \sigma_{\epsilon}^2 I), P_{data})$, is bounded by:

$$\sqrt{\frac{\alpha_{t_1}^2}{4\sigma_{t_1}^2} \mathbb{E}_{\mathbf{x} \sim P_{data}} \left[\|\mathbf{x}\|_2^2 \right] + \sum_{j=2}^N \frac{\alpha_{t_j}^2}{4\sigma_{t_j}^2} t_{j-1}^2 \frac{\epsilon_{\rm cm}^2}{\Delta \tau^2}} + \frac{1}{2\sigma_\epsilon} t_N \frac{\epsilon_{\rm cm}}{\Delta \tau}} + 2dL\sigma_\epsilon.$$

Compared to Theorem 2, the upper bound in Theorem 3 has an additional term $2dL\sigma_{\epsilon}$. This is the "bias" induced by the additional perturbation $\mathcal{N}(0, \sigma_{\epsilon}^2 I)$. To get a tighter bound, we may choose $\sigma_{\epsilon} = \sqrt{\frac{t_N \epsilon_{\rm cm}}{4dL\Delta\tau}}$, and the upper bound becomes: $\sqrt{\frac{\alpha_{t_1}^2}{4\sigma_{t_1}^2}} \mathbb{E}_{\mathbf{x}\sim P_{\rm data}} \left[\|\mathbf{x}\|_2^2 \right] + \sum_{j=2}^N \frac{\alpha_{t_j}^2}{4\sigma_{t_j}^2} t_{j-1}^2 \frac{\epsilon_{\rm cm}^2}{\Delta\tau^2} + 2\sqrt{t_N dL} \frac{\epsilon_{\rm cm}}{\Delta\tau}$.

3.3. Case Studies on Multistep Sampling

To illustrate the theoretical guarantee and understand the benefits of multistep sampling, we conduct case studies with two common forward processes. For simplicity, we assume P_{data} to have bounded support and ignore the rounding issues when selecting sampling time schedule $\{t_i\}_{i=1:N}$ from the training time partition \mathcal{T} .

Case study 1: we consider the *Variance Preserving SDE* in Song et al. (2021) with $\beta(t) = 2$ as the forward process:

$$d\mathbf{x}_t = -\mathbf{x}_t dt + \sqrt{2} d\mathbf{w}_t, \quad \mathbf{x}_0 \sim P_{\text{data}}.$$
 (7)

This is also known as the Ornstein-Uhlenbeck (OU) process and is studied by Chen et al. (2023b) in the context of score-based generative models. The forward process defined by (7) has noise schedule $\alpha_t = e^{-t}$ and $\sigma_t^2 = 1 - e^{-2t}$ and its marginal distribution is $\mathbf{x}_t \sim \mathcal{N}(e^{-t}\mathbf{x}_0, (1 - e^{-2t})I)$ conditioning on \mathbf{x}_0 . Theorem 2

guarantees that $W_2(\hat{P}_{t_N}, P_{data})$ is bounded by

$$2R\left(\frac{e^{-2t_1}R^2}{4(1-e^{-2t_1})} + \sum_{j=2}^N \frac{e^{-2t_j}t_{j-1}^2}{4(1-e^{-2t_j})} \frac{\epsilon_{\rm cm}^2}{\Delta\tau^2}\right)^{1/4} + t_N \frac{\epsilon_{\rm cm}}{\Delta\tau}.$$
(8)

In this case study, we focus on the design of the sampling time schedule based on upper bound (8). To develop a reasonable multistep sampling procedure, we make the following two practical assumptions: $\Delta \tau \ll 1$ and $\frac{\epsilon_{\rm cm}}{\Delta \tau} < R$. The condition $\Delta \tau \ll 1$ allows for selecting a small final sampling step t_N . However, as we demonstrate below, an ultra-small t_N is not beneficial. The assumption $\frac{\epsilon_{\rm cm}}{\Delta \tau} < R$ ensures that the consistency model yields a meaningful self-consistency loss. ⁵ These assumptions are used solely for deriving the sampling time schedule; our theoretical results do not depend on them.

One strategy for designing $\{t_i\}_{i=1:N}$ is to *minimize the upper bound* (8). We first establish a lower bound on (8) as a baseline. Without loss of generality, we assume $t_1 \ge 2$. (8) can be lower bounded as:

$$(8) \ge R \sqrt{\frac{\epsilon_{\rm cm}}{\Delta \tau}} \left(\sum_{j=2}^{N} \frac{t_j}{e^{2t_j} - 1} (t_{j-1} - t_j) \right)^{1/4} + t_N \frac{\epsilon_{\rm cm}}{\Delta \tau}$$
$$\ge R \sqrt{\frac{\epsilon_{\rm cm}}{\Delta \tau}} \left(\int_{t_N}^2 \frac{x \mathrm{d}x}{e^{2x} - 1} \right)^{1/4} + t_N \frac{\epsilon_{\rm cm}}{\Delta \tau},$$

where the first step is because $0 < t_j \le t_{j-1}$ and the second step is because $\frac{x}{e^{2x}-1}$ monotonically decreases. Let $c_1, c_2 > 0$ be absolute constants, s.t. $\left(\int_{c_1}^2 \frac{x dx}{e^{2x}-1}\right)^{1/4} = c_2$. Then if $t_N \ge c_1$, $(8) \ge c_1 \frac{\epsilon_{\text{cm}}}{\Delta \tau} = \Omega\left(\frac{\epsilon_{\text{cm}}}{\Delta \tau}\right)$; if $t_N < c_1$, $(8) \ge c_2 R \sqrt{\frac{\epsilon_{\text{cm}}}{\Delta \tau}} = \Omega\left(R \sqrt{\frac{\epsilon_{\text{cm}}}{\Delta \tau}}\right)$. In either case, (8) =

⁵When $\frac{\epsilon_{\rm cm}}{\Delta \tau} \ge R$, (8) = $\Omega(R)$, which is uninformative because the support of $P_{\rm data}$ is already bounded by R. This trivial scenario is not the focus of this case study.

$$\begin{split} &\Omega\big(\min\big\{\frac{\epsilon_{\rm cm}}{\Delta\tau},R\sqrt{\frac{\epsilon_{\rm cm}}{\Delta\tau}}\big\}\big). \text{ The condition } \frac{\epsilon_{\rm cm}}{\Delta\tau} < R \text{ further implies } (8) = \Omega\big(\frac{\epsilon_{\rm cm}}{\Delta\tau}\big). \text{ Given this lower bound, one heuristic is to set every term in (8) to } \tilde{\Theta}\big(\frac{\epsilon_{\rm cm}}{\Delta\tau}\big)\big) \text{ to match this baseline approximately, which requires:} \end{split}$$

$$t_i \ge \log \frac{R^3 \Delta \tau^2}{\epsilon_{\rm cm}^2}$$
, if $i = 1$; $t_i \ge \log \frac{R^2 \Delta \tau}{\epsilon_{\rm cm}}$, o.w. (9)

With this heuristic, a two-step sampling procedure shows an improvement in sample quality:

Corollary 1 (Two-step sampling with OU process). Suppose the conditions in Theorem 2 are satisfied. Suppose $\alpha_t = e^{-t}, \ \sigma_t^2 = 1 - e^{-2t}$. Then for $t_1 = \log \frac{R^3 \Delta \tau^2}{\epsilon_{\rm cm}^2}$, $t_2 = \log \frac{R^2 \Delta \tau}{\epsilon_{\rm cm}}$, we have:

$$\begin{cases} W_2(\hat{P}_0^{(1)}, P_{data}) \leq \frac{\epsilon_{\rm cm}}{\Delta \tau} \left(\log \frac{R^3 \Delta \tau^2}{\epsilon_{\rm cm}^2} + O(1) \right), \\ W_2(\hat{P}_0^{(2)}, P_{data}) \leq \frac{\epsilon_{\rm cm}}{\Delta \tau} \left(\log \frac{R^2 \Delta \tau}{\epsilon_{\rm cm}} + O\left(\sqrt{\log \frac{R^2 \Delta \tau}{\epsilon_{\rm cm}}}\right) \right). \end{cases}$$
(10)

Because $\frac{\epsilon_{\rm cm}}{\Delta \tau} < R$, the leading term is strictly reduced in the second sampling step. Furthermore, if $\epsilon_{\rm cm} \approx \Delta \tau$, $W_2(\hat{P}_0^{(2)}, P_{\rm data}) \approx \frac{2}{3}W_2(\hat{P}_0^{(1)}, P_{\rm data})$; if $\epsilon_{\rm cm} \ll \Delta \tau$, $W_2(\hat{P}_0^{(2)}, P_{\rm data}) \approx \frac{1}{2}W_2(\hat{P}_0^{(1)}, P_{\rm data})$. Due to the constraint in (9), further improvement with this heuristic is challenging. This intuition aligns with the empirical result in Luo et al. (2023). Our simulation in Section I demonstrates that the sampling strategy in Corollary 1 achieves accuracy comparable to baseline strategies while requiring significantly fewer sampling steps.

Case study 2: In the second case study, we consider the following *Variance Exploding SDE* (Song et al., 2021; Karras et al., 2022) as the forward process:

$$\mathrm{d}\mathbf{x}_t = \sqrt{2t}\mathrm{d}\mathbf{w}_t,\tag{11}$$

which is used in Song et al. (2023) and Song & Dhariwal (2024) as the forward process for consistency models. The noise schedule is $(\alpha_t, \sigma_t^2) = (1, t^2)$ and the marginal distribution of \mathbf{x}_t conditioning on \mathbf{x}_0 is: $\mathbf{x}_t \sim \mathcal{N}(\mathbf{x}_0, t^2 I)$. The upper bound in (6) is simplified to:

$$\underbrace{2R\left(\frac{1}{4t_1^2}R^2 + \sum_{j=2}^{N}\frac{1}{4t_j^2}t_{j-1}^2\frac{\epsilon_{\rm cm}^2}{\Delta\tau^2}\right)^{1/4}}_{(i)} + \underbrace{t_N\frac{\epsilon_{\rm cm}}{\Delta\tau}}_{(ii)}.$$
 (12)

This implies a trade-off in multi-step sampling with this particular forward process (11) when increasing the number of steps. Roughly speaking, (i) in (12) increases due to more terms with more steps while t_N becomes smaller and (ii) will decrease. To design the sampling time schedule strategically, similar to Case study 1, we first estimate the

lower bound of equation (12). For an arbitrary sampling strategy $t_1 \ge t_2 \ge \ldots \ge t_N$, the summation inside term (i) can be lower bounded by:

$$\sum_{j=2}^{N} \frac{t_{j-1}^2}{t_j^2} \ge \sum_{j=2}^{N} \frac{t_j(t_{j-1} - t_j)}{t_j^2} = \sum_{j=2}^{N} \frac{t_{j-1} - t_j}{t_j}$$
$$\ge \int_{t_N}^{t_1} \frac{dt}{t} = \log \frac{t_1}{t_N}.$$

A schedule defined by a geometric series approximates this lower bound well. Letting $t_j = t_N \rho^{N-j}$ with $\rho > 1$, we have $\sum_{j=2}^{N} \frac{t_{j-1}^2}{t_j^2} = \rho^2 \log_{\rho} \frac{t_1}{t_N} = \frac{\rho^2}{\log \rho} \log \frac{t_1}{t_N} \ge 2e \log \frac{t_1}{t_N}$, where equality holds when $\rho = \sqrt{e}$. Substituting this into equation (12), we obtain:

$$2R\left(\frac{R^2}{4t_1^2} + \frac{2e\epsilon_{\rm cm}^2}{4\Delta\tau^2}\left(\log t_1 + \log\frac{1}{t_N}\right)\right)^{1/4} + t_N\frac{\epsilon_{\rm cm}}{\Delta\tau}.$$

To minimize this expression, we choose $t_1 = \sqrt{\frac{1}{e} \frac{R\Delta\tau}{\epsilon_{\rm cm}}}$. As $t_N \to 0$, the second term decreases linearly, while the first term increases slowly. Therefore, a reasonable choice is $t_N = \Delta\tau$, a small constant. To summarize, the t_i 's are specified by: (ignore the rounding issue)

$$t_i = t_1 e^{\frac{1-i}{2}}, \quad i = 1, 2, \dots, N,$$
 (13)

where $t_1 = \sqrt{\frac{1}{e} \frac{R\Delta\tau}{\epsilon_{\rm cm}}}$ and $N = 1 + 2\log \frac{t_1}{\Delta\tau}$. With this choice, we obtain the following error guarantee:

Corollary 2 (Multistep sampling with the variance exploding SDE). Suppose the conditions in Theorem 2 are satisfied. Suppose $\alpha_t = 1$, $\sigma_t^2 = t^2$. For $\{t_i\}_{i=1:N}$ defined in (13), we have: $W_2(\hat{P}_0^{(N)}, P_{data}) \leq \tilde{O}(R\sqrt{\frac{\epsilon_m}{\Delta \tau}})$.

For comparison, the best error rate of (12) is $O\left(R\left(\frac{\epsilon_{\rm cm}}{\Delta\tau}\right)^{1/3}\right)$ when using a single sampling step (N = 1). By choosing the sampling schedule strategically, we see a clear improvement on the error rate when using multiple sampling steps.

To summarize, the convergence of a forward process in (1) is characterized by $\alpha_t^2 \sigma_t^{-2}$ (according to Lemma 3). The forward process (11) has a polynomial convergence rate $\alpha_t^2 \sigma_t^{2-2} = t^{-2}$ while (7) enjoys a much faster exponential rate $\alpha_t^2 \sigma_t^{-2} \approx e^{-2t}$. The exponential convergence results in a *shorter training step T*, *fewer sampling steps N*, and *better sample quality*, provided that Assumption 1 holds with the same $\epsilon_{\rm cm}$ in both cases.

4. Technical Overview

In this section, we present the high-level ideas in the proof for our main result Theorem 2 since proof for Theorem 3 shares the same main building blocks. The proof for Theorem 2 consists of three main components: **Error decomposition:** intuitively, the error comes from: (i) inaccurate consistency function $\hat{f}(\cdot, \cdot)$ and (ii) sampling from Gaussian distribution $\mathcal{N}(0, \sigma_{t_1}^2)$ instead of perturbed data distribution P_{t_1} . (i) is controlled by the consistency loss Assumption 1 and (ii) is controlled by the convergence of the forward process Lemma 3. However, the error (i) and (ii) interact with each other in the multi-step sampling. We handle this complication progressively, starting with the error decomposition in the final sampling step:

$$\begin{split} W_2(\hat{P}_0^{(N)}, P_{\text{data}}) \leq & W_2(\hat{f}(\hat{P}_{t_N}, t_N), \hat{f}(P_{t_N}, t_N)) \\ & + W_2(\hat{f}(P_{t_N}, t_N), f^{\star}(P_{t_N}, t_N)). \end{split}$$

Since the output of $\hat{f}(\cdot, \cdot)$ is *bounded*, we could simplify the first term with the TV distance, which is further upper bounded by $\operatorname{KL}(P_{t_N} \parallel \hat{P}_{t_N})$ by *Pinsker's inequality* and *data processing inequality*. The second term is solely controlled by the consistency loss $\epsilon_{\rm cm}$.

Recursion on KL $(P_{t_i} \parallel \hat{P}_{t_i})$: we analyze KL $(P_{t_N} \parallel \hat{P}_{t_N})$ via *induction*. First of all, the base case KL $(P_{t_1} \parallel \hat{P}_{t_1})$ is upper bounded using the *convergence of the forward process*; the induction step connects KL $(P_{t_i} \parallel \hat{P}_{t_i})$ and KL $(P_{t_{i+1}} \parallel \hat{P}_{t_{i+1}})$. According to the multi-step sampling, \hat{P}_{t_i} and $\hat{P}_{t_{i+1}}$ is connected by $\hat{f}(\cdot, t_i)$ and the forward SDE as

$$\hat{P}_{t_i} \xrightarrow{\hat{f}(\cdot,t_i)} \hat{P}_0^{(i)} \xrightarrow{\text{SDE}} \hat{P}_{t_{i+1}}$$

In this process, $\hat{f}(\cdot, \cdot)$ induced additional error while the forward SDE reduces it with convergence $\alpha_{t_{i+1}}^2 \sigma_{t_{i+1}}^{-2}$. This intuition is formalized by the error decomposition via *chain rule of KL divergence*:

$$\begin{aligned} & \operatorname{KL}(P_{t_{i+1}} \parallel P_{t_{i+1}}) \\ & \leq \operatorname{KL}(P_{t_i} \parallel \hat{P}_{t_i}) + \frac{\alpha_{t_{i+1}}^2}{2\sigma_{t_{i+1}}^2} \mathbb{E}_{\mathbf{x} \sim P_{t_i}} \left[\left\| f^{\star}(\mathbf{x}, t_i) - \hat{f}(\mathbf{x}, t_i) \right\|_2^2 \right] \end{aligned}$$

Another possibility is to construct the recursive formula for $W_2(\hat{P}_0^{(i)}, P_0)$. However, recursion on W_2 requires the translation from KL to W_2 that induces an R factor in each induction step. When $\{t_i\}_i$ is not carefully designed, the Rin each induction step results in an exploding upper bound easily. Meanwhile, this translation requires the data distribution to be bounded and hampers the application to more general data distributions.

Error of consistency function evaluation: another important building block in our proof is the evaluation error of consistency function, i.e. $\|\hat{f}(\mathbf{x},\tau_k) - f^*(\mathbf{x},\tau_k)\|_2$ for $\tau_k \in \mathcal{T}$. Assumption 1 controls the difference in $\hat{f}(\cdot,\cdot)$ and $f^*(\cdot,\cdot)$ indirectly by enforcing the consistency property. We connect the evaluation error and consistency loss via a

stepwise decomposition. Conditioning on $\mathbf{x}_{\tau_k} \sim P_{\tau_k}$, the PF-ODE (2) defines a deterministic trajectory:

$$\mathbf{x}_{\tau_k} \xrightarrow{\varphi(\tau_{k-1}; \cdot, \tau_k)} \mathbf{x}_{\tau_{k-1}} \cdots \xrightarrow{\varphi(\tau_1; \cdot, \tau_2)} \mathbf{x}_{\tau_1} \xrightarrow{\varphi(\tau_0; \cdot, \tau_1)} \mathbf{x}_{\tau_0}.$$

Assumption 1 guarantees that $\left\| \hat{f}(\mathbf{x}_{\tau_j}, \tau_j) - \hat{f}(\mathbf{x}_{\tau_{j-1}}, \tau_{j-1}) \right\|_2$ is small in the sense of L_2 error for each intermediate step j. We could make the following decomposition:

$$\left\| \hat{f}(\mathbf{x}_{\tau_k}, \tau_k) - f^{\star}(\mathbf{x}_{\tau_k}, \tau_k) \right\|_2 = \left\| \hat{f}(\mathbf{x}_{\tau_k}, \tau_k) - \mathbf{x}_0 \right\|_2$$
$$\leq \sum_{j=1}^k \left\| \hat{f}(\mathbf{x}_{\tau_j}, \tau_j) - \hat{f}(\mathbf{x}_{\tau_{j-1}}, \tau_{j-1}) \right\|_2.$$

The right-hand side is, roughly speaking $\leq \tau_k \frac{\epsilon_{\rm cm}}{\Delta \tau}$, We formalize this idea with *Minkowski inequality* in Lemma 2.

5. Conclusion

In this paper, we study the convergence of the consistency model multistep sampling procedure. We establish guarantees on the distance between the sample distribution and data distribution in terms of both Wasserstein distance and total variation distance. Our upper bound requires only mild assumptions on the data distribution.

Future research directions include providing lower bounds on multistep sampling and establishing end-to-end results on consistency models.

Acknowledgements

Wen Sun is supported by NSF IIS-2154711, NSF CAREER 2339395, DARPA LANCER: LeArning Network CybERagents, and a Sloan research fellowship.

Impact Statement

This paper presents work whose goal is to advance the theoretical understanding of the consistency model. Due to the theoretical nature of this paper, we do not feel the need to highlight the societal consequences.

References

- Bakry, D., Gentil, I., Ledoux, M., et al. *Analysis and geometry of Markov diffusion operators*, volume 103. Springer, 2014.
- Black, K., Janner, M., Du, Y., Kostrikov, I., and Levine, S. Training diffusion models with reinforcement learning. arXiv preprint arXiv:2305.13301, 2023.
- Chen, H., Lee, H., and Lu, J. Improved analysis of scorebased generative modeling: User-friendly bounds under

minimal smoothness assumptions. In *International Conference on Machine Learning*, pp. 4735–4763. PMLR, 2023a.

- Chen, S., Chewi, S., Li, J., Li, Y., Salim, A., and Zhang, A. R. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *International Conference on Learning Representations*, 2023b.
- Chen, S., Chewi, S., Lee, H., Li, Y., Lu, J., and Salim, A. The probability flow ode is provably fast. *Advances in Neural Information Processing Systems*, 36, 2024.
- Dou, Z., Chen, M., Wang, M., and Yang, Z. Theory of consistency diffusion models: Distribution estimation meets fast sampling. In *Forty-first International Conference on Machine Learning*, 2024. URL https: //openreview.net/forum?id=pAPykbqUHf.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. Advances in neural information processing systems, 33:6840–6851, 2020.
- Huang, D. Z., Huang, J., and Lin, Z. Convergence analysis of probability flow ode for score-based generative models. *arXiv preprint arXiv:2404.09730*, 2024.
- Janner, M., Du, Y., Tenenbaum, J. B., and Levine, S. Planning with diffusion for flexible behavior synthesis. arXiv preprint arXiv:2205.09991, 2022.
- Karras, T., Aittala, M., Aila, T., and Laine, S. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35: 26565–26577, 2022.
- Kim, D., Lai, C.-H., Liao, W.-H., Murata, N., Takida, Y., Uesaka, T., He, Y., Mitsufuji, Y., and Ermon, S. Consistency trajectory models: Learning probability flow ODE trajectory of diffusion. In *The Twelfth International Conference on Learning Representations*, 2024. URL https: //openreview.net/forum?id=ymj18feDTD.
- Kingma, D., Salimans, T., Poole, B., and Ho, J. Variational diffusion models. *Advances in neural information* processing systems, 34:21696–21707, 2021.
- Lee, H., Lu, J., and Tan, Y. Convergence of score-based generative modeling for general data distributions. In *International Conference on Algorithmic Learning Theory*, pp. 946–985. PMLR, 2023.
- Li, G., Wei, Y., Chen, Y., and Chi, Y. Towards faster nonasymptotic convergence for diffusion-based generative models. arXiv preprint arXiv:2306.09251, 2023.

- Li, G., Huang, Y., Efimov, T., Wei, Y., Chi, Y., and Chen, Y. Accelerating convergence of score-based diffusion models, provably. arXiv preprint arXiv:2403.03852, 2024a.
- Li, G., Huang, Z., and Wei, Y. Towards a mathematical theory for consistency training in diffusion models. *arXiv* preprint arXiv:2402.07802, 2024b.
- Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., and Zhu, J. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022.
- Luo, S., Tan, Y., Huang, L., Li, J., and Zhao, H. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*, 2023.
- Lyu, J., Chen, Z., and Feng, S. Convergence guarantee for consistency models. arXiv preprint arXiv:2308.11449, 2023.
- Oertell, O., Chang, J. D., Zhang, Y., Brantley, K., and Sun, W. Rl for consistency models: Faster reward guided text-to-image generation, 2024.
- Rolland, P. T. Y. Predicting in uncertain environments: methods for robust machine learning. Technical report, EPFL, 2022.
- Song, Y. and Dhariwal, P. Improved techniques for training consistency models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https: //openreview.net/forum?id=WNzy9bRDvG.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum? id=PxTIG12RRHS.
- Song, Y., Dhariwal, P., Chen, M., and Sutskever, I. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023.
- Xu, M., Yu, L., Song, Y., Shi, C., Ermon, S., and Tang, J. Geodiff: A geometric diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*, 2022.
- Xu, M., Powers, A. S., Dror, R. O., Ermon, S., and Leskovec, J. Geometric latent diffusion models for 3d molecule generation. In *International Conference on Machine Learning*, pp. 38592–38610. PMLR, 2023.
- Yang, M., Cho, K., Merchant, A., Abbeel, P., Schuurmans, D., Mordatch, I., and Cubuk, E. D. Scalable diffusion for materials discovery. *arXiv e-prints*, 2023.

Zhang, Q. and Chen, Y. Fast sampling of diffusion models with exponential integrator. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum? id=Loek7hfb46P.

A. The Lipschitz Condition of Consistency Function

In this section, we consider a simple example, showing that the Lipschitz coefficient of the ground truth consistency function can be exponential in problem parameters. As a result, the Lipschitz-based assumptions can be limited.

We consider a 1-D example where the target distribution $P_{data} = P_0$ is Gaussian mixture distribution $0.5N(-1, \sigma^2) + 0.5N(1, \sigma^2)$, where $0 < \sigma \ll 1$. Let $f(\cdot)$ be the consistency function that maps samples from P_T to samples from P_0 , where T is sufficiently large. We consider OU-process as the forward process so $P_T \approx N(0, 1)$. We will show that L, the Lipschitz constant of f is **at least exponential in** $\frac{1}{\sigma^2}$.

We first list a few supporting claims with brief explanations:

- C1: $f(\cdot)$ is continuous. This is because f is the solution to the PF-ODE, which is continuous under proper conditions;
- **C2:** f(0) = 0. This is by symmetry;
- C3: $\Pr_{X \sim N(0,\sigma^2)}[X \ge t] \le \exp\left(-\frac{t^2}{2\sigma^2}\right);$
- C4: for $0 < x \ll 1$, $\Pr_{X \sim P_T}[X \in [0, x]] \ge 0.1x$ for T properly large. This is because P_T is close to Gaussian for large T.

Here is the proof: Let $x \in (0, 1)$ be small enough s.t. $f(x) \in (0, 1)$. Then, by C1-4:

$$0.1x \leq \Pr_{X \sim P_T} [X \in [0, x]] = \Pr_{X \sim P_T} [f(X) \in [f(0), f(x)]] = \Pr_{Y \sim P_0} [Y \in [0, f(x)]]$$

$$\leq 0.5 \Pr_{Z \sim N(-1, \sigma^2)} [Z \in [0, f(x)]] + 0.5 \Pr_{Z \sim N(1, \sigma^2)} [Z \in [0, f(x)]]$$

$$\leq \Pr_{Z \sim N(0, \sigma^2)} [Z \geq 1 - f(x)] \leq \exp\left(-\frac{(1 - f(x))^2}{2\sigma^2}\right).$$

In particular, choose $x = 10 \exp\left(-\frac{1}{8\sigma^2}\right)$, we have:

$$f\left(10\exp\left(-\frac{1}{8\sigma^2}\right)\right) \ge \frac{1}{2}$$

By definition,

$$L = \sup_{x,y} \frac{|f(x) - f(y)|}{|x - y|} \ge \frac{f\left(10 \exp\left(-\frac{1}{8\sigma^2}\right)\right) - f(0)}{10 \exp\left(-\frac{1}{8\sigma^2}\right) - 0} \ge \frac{1}{20} \exp\left(\frac{1}{8\sigma^2}\right).$$

This confirms that the Lipschitz constant L must grow at least exponentially in $1/\sigma^2$ in this example.

B. Multistep Sampling

We present the multistep sampling procedure in Algorithm 1. Compared to Algorithm 1 of (Song et al., 2023), we allow different choices of noise schedule in Algorithm 1.

Algorithm 1 Multistep Consistency Sampling

1: Input: a trained consistency model $\hat{f}(\cdot, \cdot)$, noise schedule $\{(\alpha_t, \sigma_t^2)\}_{t \in [0,T]}$, sampling time schedule $\{t_i\}_{i=1:N}$, where $t_N = T$. 2: $\hat{\mathbf{x}}_{t_1}^{(1)} \sim \mathcal{N}(0, \sigma_{t_1}^2 I)$ 3: for i = 1 to N - 1 do 4: $\hat{\mathbf{x}}_0^{(i)} \leftarrow \hat{f}(\hat{\mathbf{x}}_{t_i}^{(i)}, t_i)$ 5: $\hat{\mathbf{x}}_{t_{i+1}}^{(i+1)} \sim \mathcal{N}(\alpha_{t_{i+1}} \hat{\mathbf{x}}_0^{(i)}, \sigma_{t_{i+1}}^2 I)$ 6: end for 7: Output: $\hat{\mathbf{x}}_0^{(N)}$.

C. Proof of Theorem 2

At a high level, we could decompose the W_2 error $W_2(\hat{P}_0^{(N)}, P_{\text{data}})$ into:

$$W_{2}(\hat{P}_{0}^{(N)}, P_{\text{data}}) \leq W_{2}(\hat{P}_{0}^{(N)}, \hat{f}(P_{t_{N}}, t_{N})) + W_{2}(\hat{f}(P_{t_{N}}, t_{N}), P_{\text{data}})$$

$$= \underbrace{W_{2}(\hat{f}(\hat{P}_{t_{N}}, t_{N}), \hat{f}(P_{t_{N}}, t_{N}))}_{=:\mathscr{A}_{1}} + \underbrace{W_{2}(\hat{f}(P_{t_{N}}, t_{N}), f^{\star}(P_{t_{N}}, t_{N}))}_{=:\mathscr{A}_{2}}.$$
(14)

In the error decomposition (14): the first term \mathscr{A}_1 is caused by an inaccurate noise distribution \hat{P}_{t_N} and is controlled by the KL divergence of P_{t_N} from \hat{P}_{t_N} . We use the chain rule of KL divergence to derive a recursive formula for $\mathrm{KL}(P_{t_i} \parallel \hat{P}_{t_i})$, where the initial term $\mathrm{KL}(P_{t_1} \parallel \hat{P}_{t_1})$ is bounded by the convergence of the forward diffusion process:

Lemma 1 (Decomposition of KL). Suppose $\hat{f}(\cdot, \cdot)$ satisfies Assumption 1, then for all i = 1, ..., N, we have:

$$\operatorname{KL}(P_{t_i} \parallel \hat{P}_{t_i}) \leq \frac{\alpha_{t_1}^2}{2\sigma_{t_1}^2} \mathbb{E}_{\mathbf{x} \sim P_{data}} \left[\|\mathbf{x}\|_2^2 \right] + \sum_{j=2}^i \frac{\alpha_{t_j}^2}{2\sigma_{t_j}^2} t_{j-1}^2 \frac{\epsilon_{\mathrm{cm}}^2}{\Delta \tau^2}$$

We defer the proof of Lemma 1 to Section C.1. Given this result, we can bound \mathcal{A}_1 as:

$$\begin{aligned} \mathscr{A}_{1} \leq & 2R\sqrt{\mathrm{TV}(\hat{f}(\hat{P}_{t_{N}}, t_{N}), \hat{f}(P_{t_{N}}, t_{N}))} \quad \left(\mathrm{By \,Section} \, 2.2.4 \text{ of (Rolland, 2022) and } \left\| \hat{f}(\mathbf{x}, t) \right\|_{2} \leq R \right) \\ \leq & 2R\left(\frac{1}{2} \operatorname{KL}(\hat{f}(P_{t_{N}}, t_{N}) \| \hat{f}(\hat{P}_{t_{N}}, t_{N}))\right)^{1/4} \quad \left(\mathrm{By \, Pinsker's \, inequality}\right) \\ \leq & 2R\left(\frac{1}{2} \operatorname{KL}(P_{t_{N}} \| \hat{P}_{t_{N}})\right)^{1/4} \quad \left(\mathrm{By \, data \, processing \, inequality}\right) \\ \leq & 2R\left(\frac{\alpha_{t_{1}}^{2}}{4\sigma_{t_{1}}^{2}} \mathbb{E}_{\mathbf{x} \sim P_{data}} \left[\|\mathbf{x}\|_{2}^{2} \right] + \sum_{j=2}^{N} \frac{\alpha_{t_{j}}^{2}}{4\sigma_{t_{j}}^{2}} t_{j-1}^{2} \frac{\epsilon_{cm}^{2}}{\Delta \tau^{2}} \right)^{1/4} \quad \left(\mathrm{By \, Lemma \, 1 \, with \, } i = N\right) \\ \leq & 2R\left(\frac{\alpha_{t_{1}}^{2}}{4\sigma_{t_{1}}^{2}} R^{2} + \sum_{j=2}^{N} \frac{\alpha_{t_{j}}^{2}}{4\sigma_{t_{j}}^{2}} t_{j-1}^{2} \frac{\epsilon_{cm}^{2}}{\Delta \tau^{2}} \right)^{1/4} \quad \left(\mathrm{Because \, sup}_{\mathbf{x} \in \mathrm{supp}(P_{data})} \|\mathbf{x}\|_{2} \leq R\right). \end{aligned}$$

The second term \mathscr{A}_2 is caused by the difference between the pre-trained consistency function $\hat{f}(\cdot, \cdot)$ and the ground truth $f^{\star}(\cdot, \cdot)$, which is controlled by the consistency loss ϵ_{cm} .

Lemma 2. Suppose $\hat{f}(\cdot, \cdot)$ satisfies Assumption 1 holds, then for all i = 0, 1, ..., M, we have:

(i)
$$\mathbb{E}_{\mathbf{x}\sim P_{\tau_i}} \left[\left\| \hat{f}(\mathbf{x},\tau_i) - f^{\star}(\mathbf{x},\tau_i) \right\|_2^2 \right] \le \tau_i^2 \frac{\epsilon_{\text{cm}}^2}{\Delta \tau^2};$$

(ii) $W_2(\hat{f}(P_{\tau_i},\tau_i), f^{\star}(P_{\tau_i},\tau_i)) \le \tau_i \frac{\epsilon_{\text{cm}}}{\Delta \tau}.$

We defer the proof of Lemma 2 to Section C.1. Part (ii) of Lemma 2 shows that:

$$\mathscr{A}_2 \le t_N \frac{\epsilon_{\rm cm}}{\Delta \tau}.\tag{16}$$

We finish the proof of Theorem 2 by combining (15) and (16).

C.1. Proof of Auxiliary lemmas

Proof of Lemma 1. We prove this statement via induction. At a high level, the base is proved by the convergence of the forward process Lemma 3. We show the induction step by the chain rule of KL.

We use $\mathcal{D}(\cdot; \alpha_t, \sigma_t^2)$ to denote the operator on distributions defined by a noise schedule (α_t, σ_t^2) . Specifically, given any distribution P, $\mathcal{D}(P; \alpha_t, \sigma_t^2)$ is the marginal distribution of \mathbf{x}' , where $\mathbf{x}' | \mathbf{x} \sim \mathcal{N}(\alpha_t \mathbf{x}, \sigma_t^2)$ and $\mathbf{x} \sim P$. When it is clear from the context, we use $\mathcal{D}(\cdot, t)$ as a shorthand. When i = 1, we can write $\hat{P}_{t_1} = \mathcal{N}(0, \sigma_{t_1}^2)$ with the diffusion operator and a the dirac distribution:

$$\hat{P}_{t_1} = \mathcal{D}\left(\delta_0; \alpha_{t_1}, \sigma_{t_1}^2\right),\,$$

where δ_0 is the delta distribution at 0. By definition, $P_{t_1} = \mathcal{D}(P_0; \alpha_{t_1}, \sigma_{t_1}^2)$. By Lemma 3,

$$\begin{split} \operatorname{KL}(P_{t_1} \parallel \hat{P}_{t_1}) = \operatorname{KL}(\mathcal{D}\left(P_0; \alpha_{t_1}, \sigma_{t_1}^2\right) \parallel \mathcal{D}\left(\delta_0; \alpha_{t_1}, \sigma_{t_1}^2\right)) \\ \leq & \frac{\alpha_{t_1}^2}{2\sigma_{t_1}^2} W_2^2(P_0, \delta_0) = \frac{\alpha_{t_1}^2}{2\sigma_{t_1}^2} \mathbb{E}_{\mathbf{x} \in P_{\text{data}}} \left[\|\mathbf{x}\|_2^2 \right]. \end{split}$$

Thus the statement holds for i = 1. Suppose the statement holds for i = k, i.e.

$$\operatorname{KL}(P_{t_k} \parallel \hat{P}_{t_k}) \leq \frac{\alpha_{t_1}^2}{2\sigma_{t_1}^2} \mathbb{E}_{\mathbf{x} \sim P_{\text{data}}} \left[\|\mathbf{x}\|_2^2 \right] + \sum_{j=2}^k \frac{\alpha_{t_j}^2}{2\sigma_{t_j}^2} t_{j-1}^2 \frac{\epsilon_{\text{cm}}^2}{\Delta \tau^2}.$$
(17)

We first explicitly write the sequence of random variables in the multistep inference:

$$\hat{\mathbf{x}}_{t_1}^{(1)} \to \hat{\mathbf{x}}_0^{(1)} \to \hat{\mathbf{x}}_{t_2}^{(2)} \to \hat{\mathbf{x}}_0^{(2)} \to \dots \to \hat{\mathbf{x}}_{t_N}^{(N)} \to \hat{\mathbf{x}}_0^{(N)},$$

where $\hat{\mathbf{x}}_{t_1}^{(1)} \sim \mathcal{N}(0, \sigma_{t_1}^2 I), \hat{\mathbf{x}}_0^{(i)} = \hat{f}(\hat{\mathbf{x}}_0^{(i)}, t_i), \hat{\mathbf{x}}_{t_{i+1}}^{(i+1)} \sim \mathcal{N}(\alpha_{t_{i+1}} \hat{\mathbf{x}}_0^{(i+1)}, \sigma_{t_{i+1}}^2 I)$. Similarly, we also define the following process that starts at the ground truth noise distribution P_{t_1} and evolves using the ground truth consistency function $f^*(\cdot, \cdot)$:

$$\mathbf{x}_{t_1}^{(1)} \to \mathbf{x}_0^{(1)} \to \mathbf{x}_{t_2}^{(2)} \to \mathbf{x}_0^{(2)} \to \dots \to \mathbf{x}_{t_N}^{(N)} \to \mathbf{x}_0^{(N)}$$

where $\mathbf{x}_{t_1}^{(1)} \sim P_{t_1}, \mathbf{x}_0^{(i)} = f^*(\mathbf{x}_0^{(i)}, t_i), \mathbf{x}_{t_{i+1}}^{(i+1)} \sim \mathcal{N}(\alpha_{t_{i+1}} \mathbf{x}_0^{(i)}, \sigma_{t_{i+1}}^2 I).$

By the chain rule of KL divergence, we have:

$$\begin{split} & \operatorname{KL}(\mathcal{P}\left(\mathbf{x}_{t_{k+1}}^{(k+1)}\right) \parallel \mathcal{P}\left(\hat{\mathbf{x}}_{t_{k+1}}^{(k+1)}\right)) \\ & + \underbrace{\mathbb{E}_{\mathbf{x} \sim \mathcal{P}\left(\mathbf{x}_{t_{k+1}}^{(k+1)}\right)} \left[\operatorname{KL}(\mathcal{P}\left(\mathbf{x}_{t_{k}}^{(k)} | \mathbf{x}_{t_{k+1}}^{(k+1)} = \mathbf{x}\right) \parallel \mathcal{P}\left(\hat{\mathbf{x}}_{t_{k}}^{(k)} | \hat{\mathbf{x}}_{t_{k+1}}^{(k+1)} = \mathbf{x}\right))\right]}_{\geq 0} \\ & = \operatorname{KL}(\mathcal{P}\left(\mathbf{x}_{t_{k}}^{(k)}, \mathbf{x}_{t_{k+1}}^{(k+1)}\right) \parallel \mathcal{P}\left(\hat{\mathbf{x}}_{t_{k}}^{(k)}, \hat{\mathbf{x}}_{t_{k+1}}^{(k+1)}\right)) \\ & = \operatorname{KL}(\mathcal{P}\left(\mathbf{x}_{t_{k}}^{(k)}\right) \parallel \mathcal{P}\left(\hat{\mathbf{x}}_{t_{k}}^{(k)}\right) + \mathbb{E}_{\mathbf{x} \sim \mathcal{P}\left(\mathbf{x}_{t_{k}}^{(k)}\right)} \left[\operatorname{KL}(\mathcal{P}\left(\mathbf{x}_{t_{k+1}}^{(k+1)} | \mathbf{x}_{t_{k}}^{(k)} = \mathbf{x}\right) \parallel \mathcal{P}\left(\hat{\mathbf{x}}_{t_{k+1}}^{(k+1)} | \hat{\mathbf{x}}_{t_{k}}^{(k)} = \mathbf{x}\right))\right] \end{split}$$

where we use $\mathcal{P}(\mathbf{x})$ to denote the distribution of random variable \mathbf{x} . Because KL is non-negative, we have:

$$\operatorname{KL}(\mathcal{P}\left(\mathbf{x}_{t_{k+1}}^{(k+1)}\right) \parallel \mathcal{P}\left(\hat{\mathbf{x}}_{t_{k+1}}^{(k+1)}\right))$$

$$\leq \operatorname{KL}(\mathcal{P}\left(\mathbf{x}_{t_{k}}^{(k)}\right) \parallel \mathcal{P}\left(\hat{\mathbf{x}}_{t_{k}}^{(k)}\right)) + \mathbb{E}_{\mathbf{x}\sim\mathcal{P}\left(\mathbf{x}_{t_{k}}^{(k)}\right)} \left[\operatorname{KL}(\mathcal{P}\left(\mathbf{x}_{t_{k+1}}^{(k+1)} | \mathbf{x}_{t_{k}}^{(k)} = \mathbf{x}\right) \parallel \mathcal{P}\left(\hat{\mathbf{x}}_{t_{k+1}}^{(k+1)} | \hat{\mathbf{x}}_{t_{k}}^{(k)} = \mathbf{x}\right))\right]$$

By definition, this means:

$$\begin{split} & \operatorname{KL}(P_{t_{k+1}} \parallel \hat{P}_{t_{k+1}}) \\ \leq & \operatorname{KL}(P_{t_{k}} \parallel \hat{P}_{t_{k}}) + \mathbb{E}_{\mathbf{x} \sim P_{t_{k}}} \Big[\operatorname{KL}(\mathcal{N}(\alpha_{t_{k+1}} f^{\star}(\mathbf{x}, t_{k}), \sigma_{t_{k+1}}^{2} I) \parallel \mathcal{N}(\alpha_{t_{k+1}} \hat{f}(\mathbf{x}, t_{k}), \sigma_{t_{k+1}}^{2} I)) \Big] \\ = & \operatorname{KL}(P_{t_{k}} \parallel \hat{P}_{t_{k}}) + \frac{\alpha_{t_{k+1}}^{2}}{2\sigma_{t_{k+1}}^{2}} \mathbb{E}_{\mathbf{x} \sim P_{t_{k}}} \Big[\Big\| f^{\star}(\mathbf{x}, t_{k}) - \hat{f}(\mathbf{x}, t_{k}) \Big\|_{2}^{2} \Big] \\ \leq & \operatorname{KL}(P_{t_{k}} \parallel \hat{P}_{t_{k}}) + \frac{\alpha_{t_{k+1}}^{2}}{2\sigma_{t_{k+1}}^{2}} t_{k}^{2} \frac{\epsilon_{cm}^{2}}{\Delta \tau^{2}} \quad (\text{By part (i) of Lemma 2)} \\ \leq & \frac{\alpha_{t_{1}}^{2}}{2\sigma_{t_{1}}^{2}} \mathbb{E}_{\mathbf{x} \sim P_{data}} \Big[\|\mathbf{x}\|_{2}^{2} \Big] + \sum_{j=2}^{k+1} \frac{\alpha_{t_{j}}^{2}}{2\sigma_{t_{j}}^{2}} t_{j-1}^{2} \frac{\epsilon_{cm}^{2}}{\Delta \tau^{2}} \quad (\text{By (17)}) \end{split}$$

Proof of Lemma 2. We first prove part (i) with induction on t. By the definition of $f^*(\cdot, \cdot)$ in (3),

$$f^{\star}(\mathbf{x}, 0) = \varphi(0; \mathbf{x}, 0) = \mathbf{x}, \quad \forall \mathbf{x} \in \mathbb{R}^d.$$

By Assumption 1, $\hat{f}(\mathbf{x}, 0) = \mathbf{x}$ for all \mathbf{x} . Thus

$$\mathbb{E}_{\mathbf{x}\sim P_0}\left[\left\|\hat{f}(\mathbf{x},0) - f^{\star}(\mathbf{x},0)\right\|_2^2\right] = \mathbb{E}_{\mathbf{x}\sim P_0}\left[\left\|\mathbf{x} - \mathbf{x}\right\|_2^2\right] = 0,$$

which means (i) holds for i = 0.

Suppose (i) holds for i = s, i.e.

$$\sqrt{\mathbb{E}_{\mathbf{x}\sim P_{\tau_s}}\left[\left\|\hat{f}(\mathbf{x},\tau_s) - f^{\star}(\mathbf{x},\tau_s)\right\|_2^2\right]} \le \tau_s \epsilon_{\rm cm} / \Delta \tau.$$
(18)

By the property of the PF-ODE (2),

$$\varphi(\tau_{s+1}; \mathbf{x}, \tau_s) \sim P_{\tau_{s+1}}, \quad \text{if } \mathbf{x} \sim P_{\tau_s}.$$
 (19)

When i = s + 1, we have:

$$\begin{split} &\sqrt{\mathbb{E}_{\mathbf{x}'\sim P_{\tau_{s+1}}}\left[\left\|\hat{f}(\mathbf{x}',\tau_{s+1})-f^{\star}(\mathbf{x}',\tau_{s+1})\right\|_{2}^{2}\right]} \\ =&\sqrt{\mathbb{E}_{\mathbf{x}\sim P_{\tau_{s}}}\left[\left\|\hat{f}(\varphi(\tau_{s+1};\mathbf{x},\tau_{s}),\tau_{s+1})-f^{\star}(\varphi(\tau_{s+1};\mathbf{x},\tau_{s}),\tau_{s+1})\right\|_{2}^{2}\right]} \quad (\text{By (19)}) \\ =&\sqrt{\mathbb{E}_{\mathbf{x}\sim P_{\tau_{s}}}\left[\left\|\hat{f}(\varphi(\tau_{s+1};\mathbf{x},\tau_{s}),\tau_{s+1})-f^{\star}(\mathbf{x},\tau_{s})\right\|_{2}^{2}\right]} \quad (\text{By the definition of } f^{\star}(\cdot,\cdot)) \\ \leq&\sqrt{\mathbb{E}_{\mathbf{x}\sim P_{\tau_{s}}}\left[\left\|\hat{f}(\varphi(\tau_{s+1};\mathbf{x},\tau_{s}),\tau_{s+1})-\hat{f}(\mathbf{x},\tau_{s})\right\|_{2}^{2}\right]} + \sqrt{\mathbb{E}_{\mathbf{x}\sim P_{\tau_{s}}}\left[\left\|\hat{f}(\mathbf{x},\tau_{s})-f^{\star}(\mathbf{x},\tau_{s})\right\|_{2}^{2}\right]} \\ &(\text{By Lemma 5}) \\ \leq&\epsilon_{\text{cm}}+\tau_{s}\epsilon_{\text{cm}}/\Delta\tau \quad (\text{By Assumption 1 and (18)}) \\ =&\epsilon_{\text{cm}}(1+\tau_{s}/\Delta\tau)=\tau_{s+1}\epsilon_{\text{cm}}/\Delta\tau. \end{split}$$

We complete the proof for part (i).

 $\hat{f}(\cdot,t)$ and $f^{\star}(\cdot,t)$ induce a joint distribution $\Gamma_{\mathbf{x}_{0}',\mathbf{x}_{0}}$:

$$\Pr_{(\mathbf{x}'_0,\mathbf{x}_0)\sim\Gamma_{\mathbf{x}'_0,\mathbf{x}_0}}[(\mathbf{x}'_0,\mathbf{x}_0)\in A] := \Pr_{\mathbf{x}_t\sim P_t}\left[\mathbf{x}_t\in\left\{\mathbf{x}: (\hat{f}(\mathbf{x},t), f^{\star}(\mathbf{x},t))\in A\right\}\right]$$

for any event A. With this joint distribution $\Gamma_{\mathbf{x}'_0,\mathbf{x}_0}$, the marginal distribution of \mathbf{x}'_0 is $\hat{f}(P_t,t)$ and the marginal distribution of \mathbf{x}_0 is $f^*(P_t,t)$. This means:

$$\sqrt{\mathbb{E}_{\mathbf{x}_{t}\sim P_{t}}\left[\left\|\hat{f}(\mathbf{x}_{t},t) - f^{\star}(\mathbf{x}_{t},t)\right\|_{2}^{2}\right]} = \sqrt{\mathbb{E}_{(\mathbf{x}_{0}',\mathbf{x}_{0})\sim\Gamma_{\mathbf{x}_{0}',\mathbf{x}_{0}}}\left[\left\|\mathbf{x}_{0}' - \mathbf{x}_{0}\right\|_{2}^{2}\right]} \ge W_{2}(\hat{f}(P_{t},t),f^{\star}(P_{t},t)).$$

By applying part (i), we have

$$W_2(\hat{f}(P_{\tau_i}, \tau_i), f^{\star}(P_{\tau_i}, \tau_i)) \le \tau_i \epsilon_{\rm cm} / \Delta \tau.$$

We complete the proof for part (ii).

D. Generalization to Distributions with Tail Conditions

When P_{data} satisfies some tail condition, it is sufficient to sample only from a bounded region:

Theorem 4 (W_2 error for distributions with tail condition). Suppose there exists c, C > 0 and $R \ge C$, s.t. $\Pr_{\mathbf{x} \sim P_{data}}(\|\mathbf{x}\|_2 \ge t) \le ce^{-t/C}$ for all $t \ge R$. Let $P_{data \cap \mathcal{B}(0,R)}$ be the distribution truncated from P_{data} , i.e. the conditional distribution of \mathbf{x} given $\|\mathbf{x}\|_2 \le R$ where $\mathbf{x} \sim P_{data}$. Let $\varphi_R(\cdot; \cdot, \cdot)$ be the solution to the corresponding PF-ODE and $f_R^*(\cdot, \cdot)$ be the corresponding consistency function. Let $\{P_t^R\}_{t\in[0,T]}$ be the marginal distribution of the forward process starting from $P_{data \cap \mathcal{B}(0,R)}$. If $\hat{f}(\cdot, \cdot)$ satisfies: (a) $\|\hat{f}(\mathbf{x},t)\|_2 \le R$, for all $(\mathbf{x},t) \in \mathbb{R}^d \times [0,T]$; (b) $\hat{f}(\mathbf{x},0) = \mathbf{x}$, for all \mathbf{x} ; (c) $\mathbb{E}_{\mathbf{x}_t \sim P_{\tau_i}^R} \left[\|\hat{f}(\mathbf{x}_t,\tau_i) - \hat{f}(\varphi_R(\tau_{i+1};\mathbf{x}_t,\tau_i),\tau_{i+1}) \|_2^2 \right] \le \epsilon_{cm}^2$, for all $i = 0, \ldots, M - 1$ for some $\epsilon_{cm} > 0$. Then $W_2(\hat{P}_0^{(N)}, P_{data}) \le 2R \left(\frac{\alpha_{\tau_1}^2}{4\sigma_{\tau_1}^2} R^2 + \sum_{j=2}^N \frac{\alpha_{\tau_j}^2}{4\sigma_{\tau_j}^2} t_{j-1}^2 - \frac{\epsilon_{cm}^2}{\Delta \tau^2} \right)^{1/4} + t_N \frac{\epsilon_{cm}}{\Delta \tau} + O(Re^{-\frac{R}{2C}})$.

By restricting the output of $\hat{f}(\cdot, \cdot)$ to be $\mathcal{B}(0, R)$, the Euclidean ball with radius R, we focus on learning the portion of P_{data} inside the Euclidean ball. This truncation step reduces the problem of *sampling from unbounded distribution* to *sampling from a distribution with finite support*, at the cost of introducing the additional term $O(Re^{-\frac{R}{2C}})$.

Proof. The error term can be decomposed as:

$$W_2(\hat{P}_0^{(t_N)}, P_{\text{data}}) \le W_2(\hat{P}_0^{(t_N)}, P_{\text{data}\cap\mathcal{B}(0,R)}) + W_2(P_{\text{data}\cap\mathcal{B}(0,R)}, P_{\text{data}})$$
(20)

By Theorem 2,

$$W_2(\hat{P}_0^{(t_N)}, P_{\text{data}\cap\mathcal{B}(0,R)}) \le 2R \left(\frac{\alpha_{t_1}^2}{4\sigma_{t_1}^2}R^2 + \sum_{j=2}^N \frac{\alpha_{t_j}^2}{4\sigma_{t_j}^2} t_{j-1}^2 \epsilon_{\text{cm}}^2\right)^{1/4} + t_N \epsilon_{\text{cm}}.$$

For the second term, we first note that

$$\operatorname{TV}(P_{\operatorname{data}\cap\mathcal{B}(0,R)}, P_{\operatorname{data}}) = \operatorname{Pr}_{\mathbf{x}\sim P_{\operatorname{data}}}[\|\mathbf{x}\|_2 > R] \le O(e^{-\frac{R}{C}}).$$

By Lemma 9 of (Rolland, 2022),

$$W_2(P_{\mathsf{data}\cap\mathcal{B}(0,R)}, P_{\mathsf{data}}) \le O(Re^{-\frac{R}{2C}}).$$

We finish the proof by combining these two bounds.

E. Proof of Theorem 3

At a high level, we can decompose the TV distance as follows:

$$\operatorname{TV}(\hat{P}_{0}^{(N)} * \mathcal{N}(0, \sigma_{\epsilon}^{2}I), P_{\text{data}})$$

$$\leq \operatorname{TV}(\hat{P}_{0}^{(N)} * \mathcal{N}(0, \sigma_{\epsilon}^{2}I), P_{\text{data}} * \mathcal{N}(0, \sigma_{\epsilon}^{2}I)) + \operatorname{TV}(P_{\text{data}} * \mathcal{N}(0, \sigma_{\epsilon}^{2}I), P_{\text{data}})$$

$$(21)$$

The first term can be bounded by Lemma 1 and Pinsker's inequality, which shows that the TV distance between $\hat{P}_0^{(N)}$ and P_{data} is controlled after the Gaussian perturbation. While the second term is bounded when P_{data} satisfies the smoothness assumption, which shows that the perturbation will change P_{data} only slightly. We now illustrate these ideas in detail. We

first define $\alpha_{t_{N+1}} := 1$, $\sigma_{t_{N+1}} := \sigma_{\epsilon}$, then by Pinsker's inequality and Lemma 1:

$$\begin{aligned} &\operatorname{TV}(\hat{P}_{0}^{(N)} * \mathcal{N}(0, \sigma_{\epsilon}^{2}I), P_{\text{data}} * \mathcal{N}(0, \sigma_{\epsilon}^{2}I)) \\ &\leq \sqrt{\frac{1}{2}} \operatorname{KL}(P_{\text{data}} * \mathcal{N}(0, \sigma_{\epsilon}^{2}I) \parallel \hat{P}_{0}^{(N)} * \mathcal{N}(0, \sigma_{\epsilon}^{2}I)) \\ &= \sqrt{\frac{1}{2}} \operatorname{KL}(P_{t_{N+1}} \parallel \hat{P}_{t_{N+1}}) \\ &\leq \sqrt{\frac{\alpha_{t_{1}}^{2}}{4\sigma_{t_{1}}^{2}}} \mathbb{E}_{\mathbf{x} \sim P_{\text{data}}} \left[\|\mathbf{x}\|_{2}^{2} \right] + \sum_{j=2}^{N+1} \frac{\alpha_{t_{j}}^{2}}{4\sigma_{t_{j}}^{2}} t_{j-1}^{2} \epsilon_{\text{cm}}^{2} \\ &= \sqrt{\frac{\alpha_{t_{1}}^{2}}{4\sigma_{t_{1}}^{2}}} \mathbb{E}_{\mathbf{x} \sim P_{\text{data}}} \left[\|\mathbf{x}\|_{2}^{2} \right] + \sum_{j=2}^{N} \frac{\alpha_{t_{j}}^{2}}{4\sigma_{t_{j}}^{2}} t_{j-1}^{2} \epsilon_{\text{cm}}^{2} + \frac{1}{4\sigma_{\epsilon}^{2}} t_{N}^{2} \epsilon_{\text{cm}}^{2} \\ &\leq \sqrt{\frac{\alpha_{t_{1}}^{2}}{4\sigma_{t_{1}}^{2}}} \mathbb{E}_{\mathbf{x} \sim P_{\text{data}}} \left[\|\mathbf{x}\|_{2}^{2} \right] + \sum_{j=2}^{N} \frac{\alpha_{t_{j}}^{2}}{4\sigma_{t_{j}}^{2}} t_{j-1}^{2} \epsilon_{\text{cm}}^{2} + \frac{1}{2\sigma_{\epsilon}} t_{N} \epsilon_{\text{cm}}^{2}. \end{aligned}$$

On the other hand, by Lemma 4,

$$\operatorname{TV}(P_{\text{data}} * \mathcal{N}(0, \sigma_{\epsilon}^2 I), P_{\text{data}}) \leq 2dL\sigma_{\epsilon}.$$

We complete the proof by combining these two bounds into the decomposition in (21).

F. Connection to Consistency Distillation

Our Assumption 1 assumes that the self-consistency property is satisfied approximately, which aligns with both consistency distillation (Song et al., 2023). For simplicity, we consider an OU process to be the forward process:

$$\mathrm{d}\mathbf{x}_t = -\mathbf{x}_t \mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{w}_t, \quad \mathbf{x}_0 \sim P_{\mathrm{data}}$$

Given the pre-trained score function $s(\mathbf{x}, t)$, we train a consistency model from the following ODE:

$$\frac{\mathrm{d}\mathbf{x}_t}{\mathrm{d}t} = -\mathbf{x}_t - s(\mathbf{x}_t, t), \quad \mathbf{x}_T \sim \mathcal{N}(0, (1 - e^{-2T})I).$$
(22)

We assume access to an ODE solver, which can calculate φ^s , the solution to (22), exactly. Even though this solver can be computationally expensive during the training procedure, the consistency model will still be computationally efficient during the inference time.

To avoid distribution shift, we optimize the consistency loss objective (4) using the data generated from (22), instead of that from P_t , the marginal distribution of the forward process. When optimized properly, we can find a \hat{f} , s.t.

$$\mathbb{E}_{\mathbf{x}_{\tau_{i}}\sim\varphi^{s}(\tau_{i};\mathcal{N}(0,(1-e^{-2T})I),T)}\left[\left\|\hat{f}(\mathbf{x}_{\tau_{i}},\tau_{i})-\hat{f}(\varphi(\tau_{i+1};\mathbf{x}_{\tau_{i}},\tau_{i}),\tau_{i+1})\right\|_{2}^{2}\right]$$
(23)

is small for all *i*. Using the same argument in Lemma 4, we can show that $\hat{f}(\mathcal{N}(0, (1 - e^{-2T})I), T))$ and $\varphi^s(0; \mathcal{N}(0, (1 - e^{-2T})I), T))$ are close in W_2 , this can be translated into a bound in TV using the argument in Section 3.2.

When the pre-trained score function $s(\mathbf{x}, t)$ has small L_2 error, (Huang et al., 2024) show that $\varphi^s(0; \mathcal{N}(0, (1 - e^{-2T})I), T)$ is close to P_{data} in TV. To conclude, $\hat{f}(\mathcal{N}(0, (1 - e^{-2T})I), T)$ is close to P_{data} in TV.

G. Adaptation to Non-uniform Discretization

In this paper, we adopt a uniform discretization for clarity and ease of presentation. However, our results can be extended to the non-uniform discretization setting as well. Suppose $\tau_{0:M}$ is an arbitrary discretization of the interval [0, T]. In this scenario, it is reasonable to assume that the consistency loss scales with the length of the discretization interval:

$$E_{x_{\tau_i} \sim P_{\tau_i}} \left[\left\| \hat{f}(x_{\tau_i}, \tau_i) - \hat{f}(\varphi(\tau_{i+1}; x_{\tau_i}, \tau_i), \tau_{i+1}) \right\|_2^2 \right] \le (\tau_{i+1} - \tau_i)^2 \epsilon^2.$$

Using the same argument as in the proof of Lemma 2, we can derive:

$$\sqrt{E_{x_{\tau_i} \sim P_{\tau_i}} \left[\left\| \hat{f}(x_{\tau_i}, \tau_i) - f^{\star}(x_{\tau_i}, \tau_i) \right\|_2^2 \right]} \le \sum_{s=0}^{i-1} \sqrt{E_{x_{\tau_s} \sim P_{\tau_s}} \left[\left\| \hat{f}(x_{\tau_s}, \tau_s) - \hat{f}(\varphi(\tau_{s+1}; x_{\tau_s}, \tau_s), \tau_{s+1}) \right\|_2^2 \right]},$$

which is upper bounded by: $\sum_{s=0}^{i-1} (\tau_{s+1} - \tau_s) \epsilon = \tau_i \epsilon$. The rest of the proof remains unchanged.

H. Technical Lemmas

We first present the result on the convergence of SDE, which also connects KL-divergence and W_2 : Lemma 3. Let P and Q be two distributions in \mathbb{R}^d , then

$$\mathrm{KL}(\mathcal{D}\left(P;\alpha,\sigma^{2}\right) \parallel \mathcal{D}\left(Q;\alpha,\sigma^{2}\right)) \leq \frac{\alpha^{2}}{2\sigma^{2}}W_{2}^{2}(P,Q),$$

where we use $\mathcal{D}(P; \alpha, \sigma^2)$ to denote the marginal distribution of \mathbf{x}' , with $\mathbf{x}' | \mathbf{x} \sim \mathcal{N}(\alpha_t \mathbf{x}, \sigma_t^2)$ and $\mathbf{x} \sim P$.

This result is comparable to Lemma C.4 of (Chen et al., 2023a). However, our results is self-contained and tighter.

Proof of Lemma 3. Let U and V be two random variables with joint distribution Γ , s.t. the marginal distributions of U and V are P and Q respectively. Let $X \sim \mathcal{D}(P; \alpha, \sigma^2)$ and $Y \sim \mathcal{D}(Q; \alpha, \sigma^2)$. We use $\mathcal{P}(\cdot)$ to denote the distribution of a random variable. By the chain rule of KL-divergence, we have:

$$\begin{aligned} \operatorname{KL}(\mathcal{P}(X) \parallel \mathcal{P}(Y)) &\leq \operatorname{KL}(\mathcal{P}(X) \parallel \mathcal{P}(Y)) + \mathbb{E}_{\mathbf{x} \sim \mathcal{P}(X)}[\operatorname{KL}(\mathcal{P}((U, V) \mid X = \mathbf{x}) \parallel (U, V) \mid Y = \mathbf{x})] \\ &\quad (\text{By the non-negativity of KL}) \\ &= \operatorname{KL}(\mathcal{P}(U, V) \parallel \mathcal{P}(U, V)) \\ &\quad + \mathbb{E}_{(\mathbf{u}, \mathbf{x}) \sim \mathcal{P}(U, V)}[\operatorname{KL}(\mathcal{P}(X \mid (U, V) = (\mathbf{u}, \mathbf{v}))) \parallel \mathcal{P}(Y \mid (U, V) = (\mathbf{u}, vb)))] \\ &\quad (\text{By the chain rule of KL}) \\ &= \mathbb{E}_{(\mathbf{u}, \mathbf{x}) \sim \mathcal{P}(U, V)}[\operatorname{KL}(\mathcal{P}(X \mid U = \mathbf{u})) \parallel \mathcal{P}(Y \mid V = \mathbf{v}))] \\ &\quad (X \text{ is independent of } V \text{ given } U \text{ and similar holds for } Y) \end{aligned}$$

By the definition of $\mathcal{D}(\cdot; \cdot, \cdot)$, $X|U = \mathbf{u} \sim \mathcal{N}(\alpha \mathbf{u}, \sigma^2 I)$ and $Y|V = \mathbf{v} \sim \mathcal{N}(\alpha \mathbf{v}, \sigma^2 I)$. Thus,

$$\mathrm{KL}(\mathcal{P}(X|U=\mathbf{u})) \parallel \mathcal{P}(Y|V=\mathbf{v})) = \frac{1}{2\sigma^2} \alpha^2 \|\mathbf{u} - \mathbf{v}\|_2^2$$

By (24), we further have:

$$\operatorname{KL}(\mathcal{D}(P;\alpha,\sigma^{2}) \parallel \mathcal{D}(Q;\alpha,\sigma^{2})) \leq \frac{\alpha^{2}}{2\sigma^{2}} \mathbb{E}_{(\mathbf{u},\mathbf{v})\sim\Gamma} \left[\|\mathbf{u}-\mathbf{v}\|_{2}^{2} \right]$$
(25)

By taking inf over Γ on both sides of (25), we get:

$$\mathrm{KL}(\mathcal{D}(P;\alpha,\sigma^2) \parallel \mathcal{D}(Q;\alpha,\sigma^2)) \leq \frac{\alpha^2}{2\sigma^2} W_2^2(P,Q).$$

Lemma 4 (Gaussian perturbation on a smooth distribution, a variant of Lemma 6.4 of (Lee et al., 2023)). Let P be a distribution in \mathbb{R}^d with PDF $p(\mathbf{x})$, if $\log p(\mathbf{x})$ is L-smooth, then

$$\operatorname{TV}(P, P * \mathcal{N}(0, \sigma^2 I)) \le 2dL\sigma,$$

where we use P * Q to denote the convolution of distribution P and Q.

Proof. The results follows directly from Lemma 6.4 of (Lee et al., 2023) with $\alpha_t = 1$ and $\sigma_t = \sigma$.

Lemma 5 (Triangle inequality with both L_p norm and L_2 norm). Let \mathbf{x} be a random variable in \mathbb{R}^d , and f, g be mappings from \mathbb{R}^d to \mathbb{R}^d , then

$$\mathbb{E}_{\mathbf{x}}[\|f(\mathbf{x}) + g(\mathbf{x})\|_{2}^{p}]^{1/p} \le \mathbb{E}_{\mathbf{x}}[\|f(\mathbf{x})\|_{2}^{p}]^{1/p} + \mathbb{E}_{\mathbf{x}}[\|g(\mathbf{x})\|_{2}^{p}]^{1/p}$$

Proof.

$$\begin{split} \mathbb{E}_{\mathbf{x}}[\|f(\mathbf{x}) + g(\mathbf{x})\|_{2}^{p}]^{1/p} \leq & \mathbb{E}_{\mathbf{x}}[(\|f(\mathbf{x})\|_{2} + \|g(\mathbf{x})\|_{2})^{p}]^{1/p} \quad (\text{Triangle inequality for } L_{2} \text{ norm}) \\ \leq & \mathbb{E}_{\mathbf{x}}[\|f(\mathbf{x})\|_{2}^{p}]^{1/p} + \mathbb{E}_{\mathbf{x}}[\|g(\mathbf{x})\|_{2}^{p}]^{1/p} \quad (\text{Minkowski inequality}). \end{split}$$

I. Simulation

Motivations: Consistency model has already demonstrated its power on large-scale image generation tasks (Luo et al., 2023; Song et al., 2023; Song & Dhariwal, 2024). To verify our theoretical findings, we focus on a toy example that is easier to interpret.

We first refine our upper bound in Theorem 2, where we relax our result for a cleaner presentation. We make adjustment to (15) and get:

$$\sup_{\mathbf{x},\mathbf{y}\in\operatorname{supp}(P_{data})} \|\mathbf{x}-\mathbf{y}\|_2 \left(\frac{\alpha_{t_1}^2}{2\sigma_{t_1}^2} \mathbb{E}_{\mathbf{x}\sim P_{data}} \left[\|\mathbf{x}\|_2^2\right] + \sum_{j=2}^N \frac{\alpha_{t_j}^2}{4\sigma_{t_j}^2} t_{j-1}^2 \frac{\epsilon_{\mathrm{cm}}^2}{\Delta\tau^2}\right)^{1/4} + t_N \frac{\epsilon_{\mathrm{cm}}}{\Delta\tau}.$$
(26)

Simulation setting: We consider OU process as the forward process, which is our setup in Case study 1. For simplicity, we consider a Bernoulli data distribution: $\Pr_{x \sim P_{data}}[x=0] = \Pr_{x \sim P_{data}}[x=100] = 0.5$. This data distribution ensures a close-form for the ground truth consistency function:

$$f^{\star}(x,t) := \begin{cases} 0 & \text{if } x < 50 \exp(-t) \\ 100 & \text{o.w.} \end{cases}$$

We construct a perturbed $\hat{f}(\cdot, \cdot)$ accordingly:

$$\hat{f}(x,t) := \begin{cases} 0 & \text{if } x < a_t \\ 100 & \text{o.w.} \end{cases}$$

where the sequence a_t satisfies: $\Pr_{x \sim P_t}[x < a_t] = 0.5 + 0.0001t^2$, $\forall t$. This choice of $\hat{f}(\cdot, \cdot)$ makes sure:

$$\mathbb{E}_{\mathbf{x}\sim P_t}\left[\left\|\hat{f}(\mathbf{x},t) - f^{\star}(\mathbf{x},t)\right\|_2^2\right] = t^2.$$

This means $\hat{f}(\cdot, \cdot)$ satisfies the first statement of Lemma 2 with $\frac{\epsilon_{cm}^2}{\Delta \tau^2} = 1$.

We simulate three instantiations of $\{t_i\}_{i=1}^N$ defined in (5), i.e. the sequence of time steps for our multi-step sampling defined in (5):

- **our schedule:** the two-step schedule suggested by **Case study 1**. We also calculate the upper bound in (26) for comparison;
- **baseline 1:** design the sequence of sampling time steps by evenly dividing an interval;
- **baseline 2:** start with some T and reduce it by half every step until reaching a small value.

In Figure 2, we plot the W_2 error in multi-step sampling. We present the revolution of W_2 error in a sampling time schedule on a single curve. Specifically, we plot each curve by:

$$\left(t_i, W_2(\hat{P}_0^{(i)}, P_{\text{data}})\right)$$
 i = 1, ..., N.

Because the sampling time step t_i decreases in the multi-step sampling by definition. We reverse the x-axis of the plot for presentation purposes.



Figure 2: W_2 error in multi-step sampling.

Observations: This simulation result demonstrates that:

- Our upper bound is a reasonable characterization of the performance for the designed sampling time schedule.
- The two-step sampling time schedule suggested by **Case study 1** achieves comparable performance to the best result in the baseline methods but with a much smaller number of function evaluations;
- Running too many sampling time steps may degrade the sampling quality. The error increases for both baseline methods in the last few sampling steps.