

LEARNING DIFFUSION BRIDGES ON CONSTRAINED DOMAINS

Xingchao Liu, Lemeng Wu, Mao Ye, Qiang Liu

Department of Computer Science

University of Texas at Austin

{xcliu, lmwu, my21, lqiang}@cs.utexas.edu

ABSTRACT

Diffusion models have achieved promising results on generative learning recently. However, because diffusion processes are most naturally applied on the unconstrained Euclidean space \mathbb{R}^d , key challenges arise for developing diffusion based models for learning data on constrained and structured domains. We present a simple and unified framework to achieve this that can be easily adopted to various types of domains, including product spaces of any type (be it bounded/unbounded, continuous/discrete, categorical/ordinal, or their mix). In our model, the diffusion process is driven by a drift force that is a sum of two terms: one singular force designed by *Doob's h-transform* that ensures all outcomes of the process to belong to the desirable domain, and one non-singular neural force field that is trained to make sure the outcome follows the data distribution statistically. Experiments show that our methods perform superbly on generating tabular data, images, semantic segments and 3D point clouds. Code is available at <https://github.com/gnobitab/ConstrainedDiffusionBridge>.

1 INTRODUCTION

Diffusion-based deep generative models, notably score matching with Langevin dynamics (SMLD) (Song & Ermon, 2019, 2020), denoising diffusion probabilistic models (DDPM) (Ho et al., 2020), and their variants (e.g., Song et al., 2020b,a; Kong & Ping, 2021; Song et al., 2021; Nichol & Dhariwal, 2021), have shown to achieve new state of the art results for image synthesis (Dhariwal & Nichol, 2021; Ramesh et al., 2022; Ho et al., 2022; Liu et al., 2021), audio synthesis (Chen et al., 2020; Kong et al., 2020), point cloud synthesis (Luo & Hu, 2021a,b; Zhou et al., 2021), and many other AI tasks. These methods train a deep neural network to drive as drift force a diffusion process to generate data, and are shown to outperform competitors, mainly GANs and VAEs, on stability and sample diversity (Xiao et al., 2021; Ho et al., 2020; Song et al., 2020b).

However, due to the continuous nature of diffusion processes, the standard approaches are restricted to generating unconstrained continuous data in \mathbb{R}^d . For generating data constrained on special structured domains, such as discrete, bound data or mixes of them, special techniques, e.g., dequantization (Uria et al., 2013; Ho et al., 2019) and multinomial diffusion (Hoogeboom et al., 2021; Austin et al., 2021), need to be developed case by case and the results still tend to be unsatisfying despite promising recent advances (Hoogeboom et al., 2021; Austin et al., 2021).

This work proposes a simple and unified framework for learning diffusion models on general constrained domains Ω embedded in the Euclidean space \mathbb{R}^d . The idea is to learn a continuous \mathbb{R}^d -valued diffusion process Z_t on time interval $t \in [0, T]$, with a carefully designed force field, such that the final state Z_T guarantees to 1) fall into the desirable domain Ω , and 2) follows the data distribution asymptotically. We achieve both steps by leveraging a key tool in stochastic calculus called *Doob's h-transform* (Doob, 1984), which provides formula for deriving the diffusion processes whose final states are guaranteed to fall into a specific set or equal a specific value.




Figure 1: An Ω -Bridge on discrete domain $\Omega = \{1, 2, 3, 4\}$.

Algorithm 1 Learning Diffusion Models on Constrained Domains (a Simple Example)

Input: A dataset $\mathcal{D} := \{x^{(i)}\}$ drawn from distribution Π^* on a domain $\Omega = \{e_1, e_2, \dots, e_K\}$.

Goal: Learn a diffusion model that terminates at time T to generate samples from Π^* .

Learning: Solve the optimization below with stochastic gradient descent (or other optimizers)

$$\theta^* = \arg \min_{\theta} \int_0^T \mathbb{E}_{x \sim \mathcal{D}} \left[\|f^\theta(Z_t, t) - \nabla_{Z_t} \log \omega^\Omega(x | Z_t, t)\|^2 \right] dt, \quad (1)$$

where

$$\omega^\Omega(x | z, t) = \frac{\exp\left(-\frac{\|z - x\|^2}{2(T-t)}\right)}{\sum_{e \in \Omega} \exp\left(-\frac{\|z - e\|^2}{2(T-t)}\right)}, \quad Z_t = \frac{t}{T}x + \left(1 - \frac{t}{T}\right)x_0 + \sqrt{\frac{t(T-t)}{T}}\xi, \quad (2)$$

with x drawn from the dataset \mathcal{D} , $\xi \sim \mathcal{N}(0, I)$, and x_0 any initial point.

Sampling: Generate sample Z_T from

$$dZ_t = \left[f^{\theta^*}(Z_t, t) + \nabla_{Z_t} \log \sum_{e \in \Omega} \exp\left(-\frac{\|Z_t - e\|^2}{2(T-t)}\right) \right] dt + dW_t, \quad Z_0 = x_0.$$

Remark When the domain Ω is a manifold (e.g., line or surface) in \mathbb{R}^d , simply replace the sum $\sum_{e \in \Omega}$ with the corresponding line or surface (in general Hausdorff) integration \int_{Ω} on Ω .

Our simple procedure can be applied to any domain Ω once a properly defined summation (for discrete sets) or integration (for continuous domains) can be evaluated. To give a quick overview on the practical intuition without invoking the mathematical theory, we show in Algorithm 1 a simple instance of the framework when the domain is a discrete set $\Omega = \{e_1, \dots, e_K\}$. The idea is to set up the diffusion model to have a form of

$$dZ_t = [f^\theta(Z_t, t) + \nabla_{Z_t} \psi^\Omega(Z_t, t)] dt + dW_t, \quad \psi^\Omega(z, t) := \log \sum_{e \in \Omega} \exp\left(-\frac{\|z - e\|^2}{2(T-t)}\right), \quad (3)$$

where the drift is a sum of a non-singular (e.g., bounded) term $f^\theta(z, t)$ which is a trainable neural force field with parameter θ , and a singular term $\nabla_z \psi^\Omega(z, t)$, which drives Z_t towards set Ω as a gradient ascent on $\psi^\Omega(z, t)$. The $\psi^\Omega(z, t)$ measures the closeness of z to set Ω , as the log-likelihood of a Gaussian mixture model (GMM) centered on the elements in Ω with variance $T - t$. When t approaches to the terminal time T , the variance $T - t$ of the GMM goes to zero, and the magnitude of $\nabla_z \psi^\Omega(z, t)$ grows to infinity, hence ensuring that Z_T must belong to Ω . In particular, note that

$$\nabla_z \psi^\Omega(z, t) = \sum_{e \in \Omega} \omega^\Omega(e | z, t) \frac{e - z}{T - t}, \quad \omega^\Omega(e | z, t) = \frac{\exp\left(-\frac{\|z - e\|^2}{2(T-t)}\right)}{\exp(\psi^\Omega(z, t))},$$

which increases with an $O(1/(T-t))$ rate as $t \rightarrow T$; here $\omega^\Omega(e | z, t)$ is the softmax probability measuring the relative closeness of z to the elements e in Ω (see also Eq (2)).

As we show in Section 2.3, once f^θ is non-singular in the sense of the mild condition of $\int_0^T \mathbb{E}[\|f^\theta(Z_t, t)\|^2] dt < +\infty$, the diffusion model in (3) guarantees to yield a final state Z_T that belongs to Ω , and hence provides a flexible model family on Ω . Moreover, as shown in Eq 1 in Algorithm 1, the neural field f^θ can be simply trained to approximate $\nabla \log \omega^\Omega(e | z, t)$ with e plugged as the data point that we expect to achieve when starting from z at time t . Intuitively, such fitted f^θ increases the relative probability of the observed data points and hence allows us to fit the data distribution. Empirically, diffusion models learned through Ω -bridge achieves favorable results in generating mixed discrete/continuous tabular data, point clouds on grids, categorical semantic segments and discrete CIFAR10 images.

Outline The rest of the paper is organized as follows. Section 2 introduces h -transform, which allows us to derive bridge processes that are guaranteed to enter specific sets at the terminal time, and Section 2.3 specifies the parametric diffusion models for Ω -bridges. Then, with the learnable diffusion models, Section 3 introduces the general learning framework along with the loss function.

2 BACKGROUND: DIFFUSION PROCESSES AND h -TRANSFORM

A diffusion process $Z = \{Z_t : t \in [0, T]\}$ on \mathbb{R}^d follows a stochastic differential equation of form

$$\mathbb{Q} : \quad dZ_t = b(Z_t, t)dt + \sigma(Z_t, t)dW_t, \quad (4)$$

where W_t is a Wiener process, and $\sigma : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a positive diffusion coefficient, and $b \in [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a drift function. We use \mathbb{Q} (or \mathbb{P}) to denote the path measure of stochastic processes Z , which are probability measures on the space of continuous paths. Let \mathbb{Q}_t be the marginal distribution of Z_t at time t under \mathbb{Q} .

Our framework heavily relies on the *bridge processes*, special stochastic processes that guarantee to achieve a deterministic value or fall into a given set at the final state T .

☞ For a set $\Omega \subseteq \mathbb{R}^d$, a process Z in \mathbb{R}^d with law \mathbb{Q} is called an Ω -bridge if $\mathbb{Q}(Z_T \in \Omega) = 1$.

One natural approach to constructing bridge processes is to derive the conditioned process of a general unconstrained process given that the desirable bridge constraint happens. Specifically, assume that \mathbb{Q} is the law of a general unconstrained diffusion process of form (4), and denote by $\mathbb{Q}^\Omega(\cdot) = \mathbb{Q}(\cdot \mid Z_T \in \Omega)$ the conditioned distribution given that the event of $Z_T \in \Omega$ happens. Then \mathbb{Q}^Ω is guaranteed to be an Ω -bridge by definition. Importantly, a remarkable result from Doob (Doob, 1984), now known as h -transform, shows that \mathbb{Q}^Ω is the law of a diffusion process with a properly modified drift term. Below, we introduce this results, first for the case x -bridge when $\Omega = \{x\}$ includes a single point, and then for more general sets Ω . For simplicity, we only state the formula from h -transform that are useful for us without proofs. See e.g., Oksendal (2013); Rogers & Williams (2000) for more background on h -transform.

2.1 x -BRIDGES

Let us first consider the x -bridge $\mathbb{Q}^x(\cdot) := \mathbb{Q}(\cdot \mid Z_T = x)$, the process \mathbb{Q} pinned at a deterministic terminal point $Z_T = x$. By the result from h -transform (see e.g., Oksendal (2013)), the conditioned process $\mathbb{Q}^x(\cdot) := \mathbb{Q}(\cdot \mid Z_T = x)$, if it exists, can be shown to be the law of

$$dZ_t = (b(Z_t, t) + \sigma^2(Z_t, t)\nabla_z \log q_{T|t}(x \mid Z_t)) dt + \sigma(Z_t, t)dW_t, \quad (5)$$

where $q_{T|t}(x|z)$ is the density function of the transition probability

$$\mathbb{Q}_{T|t}(dx \mid z) := \mathbb{Q}(Z_T \in dx \mid Z_t = z),$$

where dx denotes an infinitesimal volume centering around x . Compared with the diffusion process (4) of \mathbb{Q} , the main difference is that the conditioned process has an additional drift force $\sigma^2(z, t)\nabla_z \log q_{T|t}(x|z)$ which plays the role of steering Z_t towards the target $Z_T = x$; this is a singular force whose magnitude increases to infinity as $t \rightarrow T$, because $q_{T|t}(\cdot \mid z)$ is a delta measure centered at z when $t = T$.

In addition, by Bayes rule, the distribution of the initial state Z_0 should be given by

$$Z_0 \sim \mathbb{Q}_{0|T}(\cdot \mid x), \quad \mathbb{Q}_{0|T}(dz|x) \propto \mathbb{Q}_0(dz)q_{T|0}(x|z). \quad (6)$$

Example 2.1. If \mathbb{Q} is the law of $dZ_t = \sigma_t dW_t$, we have $\mathbb{Q}_{T|t}(\cdot|z) = \mathcal{N}(z, \beta_T - \beta_t)$, where $\beta_t = \int_0^t \sigma_s^2 ds$. Hence, following the formula in (5), $\mathbb{Q}^x := \mathbb{Q}(\cdot \mid Z_T = x)$ is the law of

$$dZ_t = \sigma_t^2 \frac{x - Z_t}{\beta_T - \beta_t} dt + \sigma_t dW_t, \quad (7)$$

and $Z_0 \sim \mathbb{Q}_{0|T}(dz) \propto \mathbb{Q}_0(dz)\phi(x \mid z, \beta_T - \beta_0)$, and $\phi(\cdot \mid \mu, \sigma^2)$ is the density function of $\mathcal{N}(\mu, \sigma^2)$. The process in (7) is known as a (time-scaled) Brownian bridge. Note that the drift in (7) grows to infinity in magnitude with a rate of $O(1/(\beta_T - \beta_t))$ as $t \rightarrow T$, which ensures that $Z_t = x$ with probability one.

Arbitrary initialization To make (5) the conditioned process of (4), the initial distribution must follow the Bayes rule in (6). However, thanks to the singular force $\nabla_z \log q_{T|t}(x|z)$, the process (5) can guarantee $Z_t^x = x$ from an arbitrary initialization once the process is well defined. When the initialization is different from (6), the process in (5) is no longer the conditioned process of (4), but it remains to be an x -bridge in that $Z_T = x$ is still guaranteed. To see why this is the case, assume that \mathbb{Q} is initialized from a deterministic point $Z_0 = x_0$. Then we would still have $Z_0 = x_0$ when conditioned on $Z_T \in \Omega$ by Bayes rule. This suggests that (5) starting from any deterministic initialization is the condition process of \mathbb{Q} with the same deterministic initialization, and is hence an x -bridge. As a result, (5) from any stochastic initialization is also an x -bridge because it can be viewed as the mixture of the processes with different deterministic initialization, all of which are x -bridges. See Appendix A.4 for a detailed analysis, in which it is shown that (5) with an arbitrary initialization can be viewed as the conditioned process of a special class of non-Markov processes called *reciprocal process*.

2.2 Ω -BRIDGES

More generally, for the law \mathbb{Q} of (4) and a set $\Omega \in \mathbb{R}^d$, the Ω -bridge $\mathbb{Q}^\Omega := \mathbb{Q}(\cdot | Z_T \in \Omega)$ follows

$$\mathbb{Q}^\Omega : \quad dZ_t = \eta^\Omega(Z_t, t)dt + \sigma(Z_t, t)dW_t, \quad (8)$$

with $\eta^\Omega(z, t) = b(z, t) + \sigma^2(z, t)\mathbb{E}_{x \sim \mathbb{Q}_{T|t, z, \Omega}}[\nabla_z \log q_{T|t}(x|z)]$, $Z_0 \sim \mathbb{Q}_{0|T}(\cdot | Z_T \in \Omega)$,

where drift force η^Ω is similar to that of the x -bridge in (5), except that the final state x is now randomly drawn from an Ω -truncated (or Ω -conditioned) transition probability:

$$\mathbb{Q}_{T|t, z, \Omega}(dx | z) := \mathbb{Q}(Z_T \in dx | Z_t = z, Z_T \in \Omega),$$

which is the transition probability from Z_t to Z_T , conditioned on that $Z_T \in \Omega$. In practice, its form can be derived using Bayes rule.

Example 2.2. Assume \mathbb{Q} follows $dZ_t = \sigma_t dW_t$. Then \mathbb{Q}^Ω yields the following Ω -bridge:

$$dZ_t = \eta^\Omega(Z_t, t)dt + \sigma_t dW_t, \quad \eta^\Omega(z, t) = \sigma_t^2 \mathbb{E}_{x \sim \mathcal{N}_\Omega(z, \beta_T - \beta_t)} \left[\frac{x - z}{\beta_T - \beta_t} \right], \quad (9)$$

where $\mathcal{N}_\Omega(z, \sigma^2) = \text{Law}(Z | Z \in \Omega)$ with $Z \sim \mathcal{N}(\mu, \sigma^2)$, which is an Ω -truncated Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$, whose density function is $\phi_\Omega(x) \propto \mathbb{I}(x \in \Omega)\phi(x|\mu, \sigma)$ with $\phi(x|\mu, \sigma)$ the density function of $\mathcal{N}(\mu, \sigma^2)$.

Note that it is tractable to calculate η^Ω once we can evaluate the expectation of $\mathcal{N}_\Omega(z, \beta_T - \beta_t)$. A general case is when $\Omega = I_1 \times \dots \times I_d$, for which the expectation reduces to one dimensional Gaussian integrals. See Appendix A.6 and A.7 for details and examples of η^Ω .

As in the x -bridge, we can set the initialization to be any distribution supported on the set of points that can reach Ω following \mathbb{Q} (precisely, points z_0 that satisfy $\Omega \cap \text{supp}(\mathbb{Q}_T(\cdot | Z_0 = z_0)) \neq \emptyset$) using the mixture of initialization argument.

2.3 A PARAMETRIC FAMILY OF Ω -BRIDGES

The formula in (8) only provides a fixed process for a given \mathbb{Q} . For the purpose of learning generative models, however, we need a rich family of Ω -bridges within which we can search for a best one to fit with the data distribution. It turns out we can achieve this by simply adding an extra non-singular drift force, which can be a trainable neural network, on top of the Ω -bridge in (8). Specifically, we construct the following parametric diffusion model \mathbb{P}^θ :

$$\mathbb{P}^\theta : \quad dZ_t = (\sigma(Z_t, t)f^\theta(Z_t, t) + \eta^\Omega(Z_t, t))dt + \sigma(Z_t, t)dW_t, \quad Z_0 \sim \mathbb{P}_0^\theta, \quad (10)$$

where $f^\theta(z, t)$ is a neural network with input (z, t) and parameter θ , which will be trained based on the empirical observations. Adding the neural drift $\sigma(Z_t, t)f^\theta(Z_t, t)$ term does not break the Ω -bridge condition, once it satisfies a very mild regularization condition:

Proposition 2.3. For any \mathbb{Q}^Ω following $dZ_t = \eta^\Omega(Z_t, t)dt + \sigma(Z_t, t)dW_t$ that is an Ω -bridge, the \mathbb{P}^θ in (10) is also an Ω -bridge if $\mathbb{E}_{Z \sim \mathbb{Q}^\Omega}[\int_0^T \|f^\theta(Z_t, t)\|_2^2 dt] < +\infty$ and $\mathcal{KL}(\mathbb{Q}_0^\Omega \| \mathbb{P}_0^\theta) < +\infty$.

The condition on f^θ is very mild, and it is satisfied if f^θ is bounded. Moreover, it can easily hold even when f_θ is not bounded. For example, assuming that $f_\theta(x) \leq a \|x\|^\beta + b$, which holds for ReLU network with $\beta = 1$, we just need to require that the underlying process has a bounded moment $E_{Z \sim Q^\Omega} [\int_0^T \|Z_t\|^{2\beta} dt] < +\infty$, which is a typical regularity condition to expect.

3 LEARNING Ω -BRIDGE MODELS

Let $\{x^{(i)}\}_{i=1}^n$ be an i.i.d. sample from an unknown distribution Π^* on a domain $\Omega \subseteq \mathbb{R}^d$. Our goal is to learn the parameter θ for the Ω -bridge model \mathbb{P}^θ in (10) such that the terminal distribution $Z_T \sim \mathbb{P}_T^\theta$ matches the data $X \sim \Pi^*$. We should distinguish \mathbb{P}^θ , which is the trainable generative model, and \mathbb{Q} , which is a fixed ‘‘baseline process’’ that helps us to derive methods for constructing and learning the model. \mathbb{Q} can be the simple Brownian motion in Example 2.1 and 2.2.

As the case of other diffusion models, \mathbb{P}^θ can be viewed as a model with an infinite dimensional latent variable of the intermediate trajectories of Z . Hence, a canonical learning approach is expectation maximization (EM), which alternates between

- 1) E-step: estimating the posterior $\mathbb{P}^{\theta,x} := \mathbb{P}^\theta(Z | Z_T = x)$ of the latent trajectories Z given the observation $Z_T = x$;
- 2) M-step: estimating the parameter θ with Z imputed from $\mathbb{P}^{\theta,x}$.

A key challenge, however, is that the posterior distribution $\mathbb{P}^{\theta,x}$ is difficult to calculate due to the presence of neural force field in \mathbb{P}^θ (as the h -transform formula would have no closed form), and it need to be iteratively updated as θ changes. Following DDPM (Ho et al., 2020), we consider a simpler approach that replaces the posterior $\mathbb{P}^{\theta,x}$ with an arbitrary x -bridge, denoted by \mathbb{Q}^x . This yields a simplified EM algorithm without the expensive posterior inference in the E-step. A natural choice is the conditioned process $\mathbb{Q}^x := \mathbb{Q}(Z | Z_T = x)$, but the method works for a general x -bridge.

Specifically, let $\mathbb{Q}^{\Pi^*}(\cdot) = \int \mathbb{Q}^x(\cdot) \Pi^*(dx)$ be the mixture of the x -bridges whose end point x is randomly drawn from the data distribution $x \sim \Pi^*$. The trajectories from \mathbb{Q}^{Π^*} can be generated in the following ‘‘backward’’ way: first drawing a data point $x \sim \Pi^*$, and then $Z \sim \mathbb{Q}^x$ conditioned on the end point x . Obviously, by construction, the terminal distribution of \mathbb{Q}^{Π^*} equals Π^* , that is, $\mathbb{Q}_T^{\Pi^*} = \Pi^*$. Then, the model \mathbb{P}^θ can be estimated by fitting data drawn from \mathbb{Q}^{Π^*} using maximum likelihood estimation:

$$\min_{\theta} \left\{ \mathcal{L}(\theta) := \mathcal{KL}(\mathbb{Q}^{\Pi^*} \parallel \mathbb{P}^\theta) \right\}. \quad (11)$$

The classical (variational) EM would alternatively update θ (M-step) and \mathbb{Q}^x (E-step) to make $\mathbb{Q}^x \approx \mathbb{P}^{\theta,x}$. Why is it OK to simply drop the E-step? At the high level, it is the benefit from using universal approximators like deep neural networks: if the model space of \mathbb{P}^θ is sufficiently rich, by minimizing the KL divergence in (11), \mathbb{P}^θ can approximate the given \mathbb{Q}^{Π^*} well enough (in a way that is made precise in the Appendix A) such that their terminal distributions are close: $\mathbb{P}_T^\theta \approx \mathbb{Q}_T^{\Pi^*} = \Pi^*$.


☞ *Learning latent variable models require no E-step if the model space is sufficiently rich.*

We should see that in this case the latent variables Z in the learned model \mathbb{P}^θ is dictated by the choice of the imputation distribution \mathbb{Q} since we have $\mathbb{P}^{\theta,x} = \mathbb{Q}^x$ when the KL divergence in (11) is fully minimized to zero; EM also achieves $\mathbb{P}^{\theta,x} = \mathbb{Q}^x$ but has the imputation distribution \mathbb{Q}^x determined by the model \mathbb{P}^θ , not the other way.

Loss Function In its general form, the x -bridge \mathbb{Q}^x that we use can be a non-Markov diffusion process

$$\mathbb{Q}^x : \quad dZ_t = \eta^x(Z, t)dt + \sigma(Z_t, t)dW_t, \quad Z_0 \sim \mu^x, \quad (12)$$

which has the same diffusion coefficient $\sigma(Z_t, t)$ as \mathbb{P}^θ in (10), and any x -dependent η^x and initialization μ^x once the x -bridge condition is ensured. As the general framework, we assume that η^x can depend on the whole trajectory Z .



Algorithm 2 Learning Ω -Bridge Diffusion Models

Input: A dataset $\mathcal{D} := \{x^{(i)}\}$ drawn from distribution Π^* on a domain Ω .

Setup: Specify an x -bridge \mathbb{Q}^x and an Ω -bridge \mathbb{Q}^Ω

$$\mathbb{Q}^x : dZ_t = \eta^x(\mathbf{Z}, t)dt + \sigma(Z_t, t)dW_t, \quad \mathbb{Q}^\Omega : dZ_t = \eta^\Omega(Z_t, t)dt + \sigma(Z_t, t)dW_t,$$

Specify the generative model \mathbb{P}^θ based on \mathbb{Q}^Ω and a neural network f^θ :

$$\mathbb{P}^\theta : dZ_t = (\sigma(Z_t, t)f^\theta(Z_t, t) + \eta^\Omega(Z_t, t))dt + \sigma(Z_t, t)dW_t, \quad Z_0 \sim \mathbb{P}_0^\theta.$$

Default: let \mathbb{Q} be the law of $dZ_t = \sigma_t dW_t$ and derive the bridges by h -transform as $\mathbb{Q}^x = \mathbb{Q}(\cdot | Z_T = x)$ in Eq (7) and $\mathbb{Q}^\Omega = \mathbb{Q}(\cdot | Z_T \in \Omega)$ in Eq (9).

Training: Estimating θ by minimizing the loss function (13) using any off-the-shelf optimizer.

Sampling: Generate sample Z_T from \mathbb{P}^θ with the trained parameter θ .

Using Girsanov theorem (e.g., Oksendal, 2013), with \mathbb{P}^θ in (10) and \mathbb{Q}^x in (12), the KL divergence in (11) can be shown to equal to

$$\mathcal{L}(\theta) = \mathbb{E}_{Z \sim \mathbb{Q}^x} \left[\underbrace{-\log p_0^\theta(Z_0)}_{\text{MLE of initial dist.}} + \frac{1}{2} \int_0^T \underbrace{\|\sigma^{-1}(Z_t, t)(s^\theta(Z_t, t) - \eta^x(\mathbf{Z}, t))\|^2}_{\text{score matching}} dt \right] + \text{const}, \quad (13)$$

where we write s^θ as the overall drift force of \mathbb{P}^θ in (10), that is,

$$s^\theta(z, t) = \sigma(z, t)f^\theta(z, t) + \eta^\Omega(z, t),$$

and p_0^θ is the probability density function (PDF) of the initial distribution \mathbb{P}_0^θ . Therefore, $\mathcal{L}(\theta)$ is a sum of the negative log-likelihood of the initial distribution that encourages $\mathbb{P}_0^\theta \approx \mathbb{Q}_0^{\Pi^*}$, and a least squares loss between s^θ and η^x . In practice, we simply fix the initial distribution \mathbb{P}_0^θ to be a delta measure on a fixed point (say $x_0 = 0$), so we only need to train the drift function f^θ . Algorithm 1 shows a simple instance of the framework when the baseline process \mathbb{Q} is the standard Brownian motion $dZ_t = dW_t$, $\mathbb{Q}^x = \mathbb{Q}(\cdot | Z_T = x)$ and $\sigma(z, t) = 1$. Note that the least squares term in (13) can be viewed as enforcing $f^\theta \approx \sigma^{-1}(\eta^x - \eta^\Omega)$, which reduces to $f^\theta \approx \nabla \log \omega^\Omega$ in the case of Algorithm 1.

Related Works Bridge processes provide a simple and flexible approach to learning diffusion generative models, which was explored in Peluchetti (2021); Ye et al. (2022); Wu et al. (2022); De Bortoli et al. (2021). Heng et al. (2021) investigates the orthogonal problem of simulating from the bridge \mathbb{Q}^x for a given \mathbb{Q} . In comparison, our method learns diffusion models on general domains Ω on which an Ω -bridge can be derived (using h -transform or any other method), and hence provides a highly flexible framework for learning with structured data (including discrete, continuous, and their mixes). This distinguishes it with existing approaches that are designed for special types of data (e.g., Ho et al. (2020); Hoogeboom et al. (2021); Austin et al. (2021); Li et al. (2022); Dieleman et al. (2022) for discrete data). De Bortoli et al. (2022) discusses how to learn score-based generative models on general Riemannian manifolds. Another highly related work is Ye et al. (2022), which proposes to learn first hitting diffusion models for generating data on both discrete sets and spheres. The advantage of our approach is that it is simpler and easier to derive for more complex types of domains.

4 EXPERIMENTS

We evaluate our algorithms for generating mixed-typed tabular data, grid-valued point clouds, categorical semantic segmentation maps, discrete CIFAR10 images. We observe that Ω -bridge provides a particularly attractive and superb approach to generating data from various constrained domains.

Algorithm Overview For all experiments, we use Algorithm 2 with the default choice of \mathbb{Q}^x in (7) and \mathbb{Q}^Ω in (9). The specific form of η^Ω is derived based on the specific choice of the domain Ω . By default, we set the initialization $Z_0 = 0$ and the optimizer Adam.

	Logistic (\uparrow)	AdaBoost (\uparrow)	MLP (\uparrow)
Real Training Data	0.877 \pm 0.021	0.912 \pm 0.013	0.897 \pm 0.012
TVAE (Xu et al., 2019)	0.825 \pm 0.012	0.876 \pm 0.005	0.845 \pm 0.008
CTGAN (Xu et al., 2019)	0.649 \pm 0.014	0.841 \pm 0.021	0.843 \pm 0.016
CopulaGAN (Patki et al., 2016)	0.683 \pm 0.015	0.859 \pm 0.004	0.853 \pm 0.009
Mixed-Bridge	0.868\pm0.010	0.884\pm0.005	0.877\pm0.006

Table 1: Classification accuracy on the Adult Income dataset with different classifiers when trained with data synthesized by generative models. Real Training Data shows the upper bound of the metrics.


4.1 GENERATING MIXED-TYPE TABULAR DATA

Learning to generate tabular data is challenging, because tabular data usually contains a mixture of discrete and continuous attributes (Xu et al., 2019; Park et al., 2018). Unlike carefully designing special GANs as in previous works (Xu et al., 2019; Srivastava et al., 2017), Ω -bridge can be seamlessly applied to mixed-typed tabular data generation without any further modification. In contrast, diffusion processes that solely work on discrete domain (Austin et al., 2021; Hoogeboom et al., 2021) cannot be applied to this task.

In this experiment, we use the *Adult Income* dataset (Kohavi, 1996), which contains 30,162 training samples. The data points are described by a series of attributes, including continuous (age, capital-gain, etc.) and discrete (sex, race, etc.). We compare with conditional tabular GAN (Xu et al., 2019) (CTGAN), CopulaGAN (Patki et al., 2016), and Table VAE (Xu et al., 2019) (TVAE), which are state-of-the-art GAN-based and VAE-based generative models for mixed-typed tabular data. Following previous works (Xu et al., 2019; Patki et al., 2016), we measure the classification accuracy on the real data of logistic regression, AdaBoost classifier and MLP classifier when trained on the generated data.

In the Ω -bridge model, we set $\sigma_t = 3 \exp(-3t)$ and f^θ a 3-layer MLP. In this case, $\Omega = I_1 \times \dots \times I_{15}$, where I_1 to I_9 are discrete domains and I_{10} to I_{15} are non-negative continuous domains. For discrete domains, $I = \{e_1, \dots, e_d\}$, we have, $\eta^I(z, t) = \sigma_t^2 \nabla_z \log \sum_{e \in I} \exp\left(-\frac{\|z-e\|^2}{2(\beta_T - \beta_t)}\right)$; for non-negative continuous domains, $I = [0, +\infty)$, derivation shows $\eta^I(z, t) = \sigma_t^2 \nabla_z \log\left(F\left(\frac{z}{\sqrt{\beta_T - \beta_t}}\right)\right)$, where F is the standard Gaussian CDF. Finally, we have $\eta^\Omega(z, t) = \sum_{i=1}^{15} \eta^{I_i}(z_i, t)$ for the whole domain Ω . We set the number of diffusion steps to $K = 2000$. Results are shown in Table 1.

Result All the three different classifiers yield the highest accuracy when trained on the data generated by our method, referred to as *Mixed-Bridge* in this case. The result reflects that the data generated by Mixed-Bridge is closer to the real distribution than the baseline methods.



Method	MMD \downarrow	COV \uparrow	1-NNA \downarrow
PCD (Luo & Hu, 2021a)	13.37	46.60	58.94
\mathbb{R}^d -Bridge	13.30	46.52	59.32
Grid-Bridge	12.85	47.78	56.25

Figure 2 & Table 2: The point clouds (upper row) generated by different methods and meshes reconstructed from them (lower row). Grid-Bridge obtains more uniform points and hence better mesh thanks to the integer constraints. Numbers in the table are multiplied by 10^3 .

4.2 GENERATING INTEGER-VALUED POINT CLOUDS

A feature of point clouds in 3D objects in graphics is that they tend to distribute evenly, especially if they are discretized from a mesh. This aspect is omitted in most existing works on point cloud generation. As a result they tend to generate non-uniform points that are unsuitable for real applications, which often involve converting back to meshes with procedures like Ball-Pivoting (Bernardini




Figure 3: Results on generating categorical segmentation maps. Each pixel here an one-hot vector. Each dimension of the Ω -bridge starts from a deterministic and evolve through a stochastic trajectory to converge to either 0 or 1. The generated samples have similar visual quality to the training data.

Methods	ELBO (\downarrow)	IWBO (\downarrow)
Uniform Dequantization (Uria et al., 2013)	1.010	0.930
Variational Dequantization (Ho et al., 2019)	0.334	0.315
Argmax Flow (Softplus thres.) (Hooeboom et al., 2021)	0.303	0.290
Argmax Flow (Gumbel distr.) (Hooeboom et al., 2021)	0.365	0.341
Argmax Flow (Gumbel thres.) (Hooeboom et al., 2021)	0.307	0.287
Multinomial Diffusion (Hooeboom et al., 2021)	0.305	-
Cat.-Bridge (Constant Noise)	0.844	0.707
Cat.-Bridge (Noise Decay A)	0.276	0.232
Cat.-Bridge (Noise Decay B)	0.301	0.285
Cat.-Bridge (Noise Decay C)	0.363	0.302

Table 3: Results on the CityScapes dataset. Cat. refers to ‘Categorical’.

et al., 1999). We apply our method to generate point clouds that constrained on a integer grid which we show yields much more uniformly distributed points. To the best of our knowledge, we are the first work on integer-valued 3D point cloud generation.

A point cloud is a set of points $\{x_i\}_{i=1}^m$, $x_i \in \mathbb{R}^3$ in the 3D space, where m refers to the number of points. We apply two variants of our method: \mathbb{R}^d -Bridge and Grid-Bridge. \mathbb{R}^d -Bridge generates points in the continuous 3D space, i.e., $\Omega_{\mathbb{R}} = \mathbb{R}^{3m}$. Grid-Bridge generate points that on integer grids, $\Omega_{\text{Grid}} = \{1, \dots, 128\}^{3m}$. We fix the diffusion coefficient $\sigma_t = 1$. The number of diffusion steps K is set to 1000. We test our method on ShapeNet (Chang et al., 2015) chair models, and compare it with Point Cloud Diffusion (PCD) (Luo & Hu, 2021a), a state-of-the-art continuous diffusion-based generative model for point clouds. The neural network f^θ in our methods are the same as that of PCD for fair comparison. Qualitative results and quantitative results are shown in Figure 2 and Table 2. As common practice (Luo & Hu, 2021a,b), we measure minimum matching distance (MMD), coverage score (COV) and 1-NN accuracy (1-NNA) using Chamfer Distance (CD) with the test dataset.

Result Both \mathbb{R}^d -Bridge and Grid-Bridge get better MMD, COV, and 1-NNA than PCD. Moreover, by constraining the domain of interest to the integer grids, Grid-Bridge yields even better performance than \mathbb{R}^d -Bridge. In Figure 2, since the point clouds generated by Grid-Bridge are limited to integer grids, the reconstructed meshes from Ball-Pivoting clearly have higher quality than \mathbb{R}^d -Bridge and PCD.

4.3 GENERATING SEMANTIC SEGMENTATION MAPS ON CITYSCAPES

We consider unconditionally generating categorical semantic segmentation maps. We represent each pixels as a one-hot categorical vector. Hence the data domain is $\Omega = \{e_1, \dots, e_c\}^{h \times w}$, where c is the number of classes and e_i is the i -th c -dimensional one-hot vector, and h, w represent the height and width of the image. In CityScapes (Cordts et al., 2016), $h = 32, w = 64, c = 8$. In this experiment, we test different schedule of the diffusion coefficient σ_t , including (*Constant Noise*): $\sigma_t = 1$; (*Noise Decay A*): $\sigma_t = a \exp(-bt)$; (*Noise Decay B*): $\sigma_t = a(1 - t)$; (*Noise Decay C*) $\sigma_t = a(1 - \exp(-b(1 - t)))$. Here a and b are hyper-parameters. The number of diffusion steps K is set to 500. We measure the negative log-likelihood (NLL) of the test set using the learned models. The NLL (bits-per-dimension) is estimated with evidence lower bound (ELBO) and importance weighted bound (IWBO) (Burda et al., 2016), respectively, as in (Hooeboom et al., 2021). We compare Ω -Bridge with a state-of-the-art categorical diffusion algorithm, Argmax Flow (and Multinomial Diffusion) (Hooeboom et al., 2021), and the traditional methods, uniform dequantization (Uria

Methods	IS (\uparrow)	FID (\downarrow)	NLL (\downarrow)
Discrete			
D3PM uniform L_{vb} (Austin et al., 2021)	5.99	51.27	5.08
D3PM absorbing L_{vb} (Austin et al., 2021)	6.26	41.28	4.83
D3PM Gauss L_{vb} (Austin et al., 2021)	7.75	15.30	3.966
D3PM Gauss $L_{\lambda=0.001}$ (Austin et al., 2021)	8.54	8.34	3.975
D3PM Gauss + logistic $L_{\lambda=0.001}$	8.56	7.34	3.435
Integer-Bridge (Init. A)	8.77	6.77	3.46
Integer-Bridge (Init. B)	8.68	6.91	3.35
Integer-Bridge (Init. C)	8.72	6.94	3.40

Table 4: Discrete CIFAR10 Image Generation

et al., 2013) and variational dequantization (Ho et al., 2019). The numerical results of the baselines are directly adopted from (Hoogeboom et al., 2021), and experiment configuration is kept the same for fair comparison. The neural network f^θ is the same as (Hoogeboom et al., 2021). The results are shown in Figure 3 and Table 3. Our Ω -bridge is named *Categorical-Bridge (Cat.-Bridge)* in this experiment.

Result We observe that all the four kinds of Cat.-Bridge can successfully generate categorical semantic segments, and different noise schedules result in different empirical performance. Among the four variants of Cat.-Bridge, Cat.-Bridge with Noise Decay A yields the best ELBO and IWBO, surpassing all the other algorithms in comparison.

4.4 GENERATING DISCRETE CIFAR10 IMAGES

In this experiment, we apply three types of bridges. All of these bridges use the same output domain $\Omega = \{0, \dots, 255\}^{h \times w \times c}$, where h, w, c are the height, width and number of channels of the images, respectively. We set $\sigma_t = 3 \exp(-3t)$. We consider different initial distributions: (Init. A) $Z_0 = 128$; (Init. B) $Z_0 = \hat{\mu}_0$, (Init. C) $Z_0 \sim \mathcal{N}(\hat{\mu}_0, \hat{\sigma}_0)$, where $\hat{\mu}_0$ and $\hat{\sigma}_0$ are the empirical mean and variance of pixels in the CIFAR10 training set. The number of diffusion steps K is set to 1000. We compare with the variants of a state-of-the-art discrete diffusion model, D3PM (Austin et al., 2021). For fair comparison, we use the DDPM backbone (Ho et al., 2020) as the neural drift f^θ in our method, similar to D3PM. We report the Inception Score (IS) Salimans et al. (2016), Fréchet Inception Distance (FID) Heusel et al. (2017) and negative log-likelihood (NLL) of the test dataset. We call our method *Integer-Bridge* in this case. The results are shown in Table 4 and Figure 4.

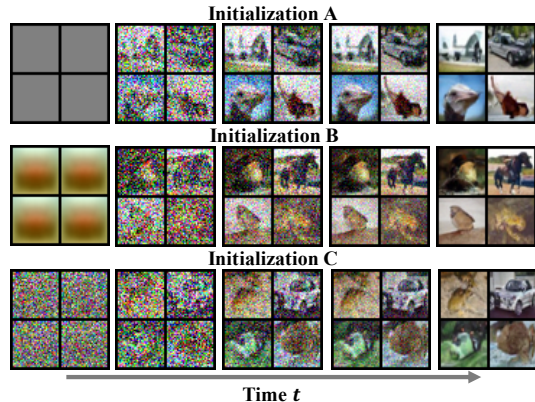


Figure 4: Integer-bridges can generate high-quality discrete samples with different initial distribution.

Result In Table 4, Integer-Bridge with Initialization A,B,C can all get lower FIDs (≤ 7) than the variants of D3PM. Among the three kinds of Integer-bridges, Integer-Bridge (Init. B) obtains the lowest NLL (3.35). It also beats D3PM Gauss + logistic (3.435) on NLL, which has the best NLL in the variants of D3PM.

5 CONCLUSION AND LIMITATIONS

We present a framework for learning diffusion generative models on constrained data domains. It leaves a number of directions for further explorations and improvement. For example, the practical impact of the choices of the bridges \mathbb{Q} , in terms of initialization, dynamics, and noise schedule, are still not well understood and need more systematical studies. Besides, our current method is limited to Ω that are factorizable and integrable. Moreover, application of Ω -bridge to many other practical fields also needs investigation in the future.

ACKNOWLEDGEMENTS

This research is supported by NSF CAREER1846421, SenSE2037267, EAGER-2041327, Office of Navy Research, and NSF AI Institute for Foundations of Machine Learning (IFML).

REFERENCES

- Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021.
- Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Cláudio Silva, and Gabriel Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE transactions on visualization and computer graphics*, 5(4):349–359, 1999.
- Yuri Burda, Roger B Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. In *ICLR (Poster)*, 2016.
- Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan. Wavegrad: Estimating gradients for waveform generation. In *International Conference on Learning Representations*, 2020.
- Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.
- Valentin De Bortoli, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. Riemannian score-based generative modeling. *arXiv preprint arXiv:2202.02763*, 2022.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34, 2021.
- Sander Dieleman, Laurent Sartran, Arman Roshannai, Nikolay Savinov, Yaroslav Ganin, Pierre H Richemond, Arnaud Doucet, Robin Strudel, Chris Dyer, Conor Durkan, et al. Continuous diffusion for categorical data. *arXiv preprint arXiv:2211.15089*, 2022.
- Joseph L Doob. *Classical potential theory and its probabilistic counterpart*, volume 549. Springer, 1984.
- Jeremy Heng, Valentin De Bortoli, Arnaud Doucet, and James Thornton. Simulating diffusion bridges with score matching. *arXiv preprint arXiv:2111.07243*, 2021.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- Jonathan Ho, Xi Chen, Aravind Srinivas, Yan Duan, and Pieter Abbeel. Flow++: Improving flow-based generative models with variational dequantization and architecture design. In *International Conference on Machine Learning*, pp. 2722–2730. PMLR, 2019.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.

- Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47):1–33, 2022.
- Emiel Hooeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in Neural Information Processing Systems*, 34, 2021.
- Ron Kohavi. Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid. In *Kdd*, volume 96, pp. 202–207, 1996.
- Zhifeng Kong and Wei Ping. On fast sampling of diffusion probabilistic models. In *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*, 2021.
- Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*, 2020.
- Antoine Lejay. The girsanov theorem without (so much) stochastic analysis. In *Séminaire de Probabilités XLIX*, pp. 329–361. Springer, 2018.
- Christian Léonard, Sylvie Røelly, and Jean-Claude Zambrini. Reciprocal processes. a measure-theoretical point of view. *Probability Surveys*, 11:237–269, 2014.
- Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. *arXiv preprint arXiv:2205.14217*, 2022.
- Xingchao Liu, Xin Tong, and Qiang Liu. Sampling with trustworthy constraints: A variational gradient framework. *Advances in Neural Information Processing Systems*, 34:23557–23568, 2021.
- Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2837–2845, 2021a.
- Shitong Luo and Wei Hu. Score-based point cloud denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4583–4592, 2021b.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pp. 8162–8171. PMLR, 2021.
- Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- Noseong Park, Mahmoud Mohammadi, Kshitij Gorde, Sushil Jajodia, Hongkyu Park, and Youngmin Kim. Data synthesis based on generative adversarial networks. *Proceedings of the VLDB Endowment*, 11(10), 2018.
- N. Patki, R. Wedge, and K. Veeramachaneni. The synthetic data vault. In *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 399–410, Oct 2016. doi: 10.1109/DSAA.2016.49.
- Stefano Peluchetti. Non-denoising forward-time diffusions. 2021.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- L Chris G Rogers and David Williams. *Diffusions, Markov processes and martingales: Volume 2, Itô calculus*, volume 2. Cambridge university press, 2000.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2020a.

- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019.
- Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2020b.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems*, 34, 2021.
- Akash Srivastava, Lazar Valkov, Chris Russell, Michael U Gutmann, and Charles Sutton. Veegan: Reducing mode collapse in gans using implicit variational learning. *Advances in neural information processing systems*, 30, 2017.
- Benigno Uribe, Iain Murray, and Hugo Larochelle. Rnade: The real-valued neural autoregressive density-estimator. *Advances in Neural Information Processing Systems*, 26, 2013.
- Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- Lemeng Wu, Chengyue Gong, Xingchao Liu, Mao Ye, and Qiang Liu. Diffusion-based molecule generation with informative prior bridges. *arXiv preprint arXiv:2209.00865*, 2022.
- Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion gans. *arXiv preprint arXiv:2112.07804*, 2021.
- Lei Xu, Maria Skoularidou, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. Modeling tabular data using conditional gan. *Advances in Neural Information Processing Systems*, 32, 2019.
- Mao Ye, Lemeng Wu, and Qiang Liu. First hitting diffusion models. *arXiv preprint arXiv:2209.01170*, 2022.
- Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5826–5835, 2021.

Roadmap The appendix is structured as follows:

- Appendix A provides the theoretical analysis and derivation of diffusion bridges. In particular, Appendix A.1 shows the derivation of the main training loss; Appendix A.2 derives the drift term η^{Π^*} of \mathbb{Q}^{Π^*} ; Appendix A.3 proves that we can actually use a Markov model \mathbb{P}^θ to match all time-marginals with \mathbb{Q}^{Π^*} ; Appendix A.4 explains why we can use arbitrary initialization when constructing bridge processes and discusses the reciprocal structure of \mathbb{Q}^{Π^*} ; Appendix A.5 provides analysis on the time-discretization error and statistical error of the practical discretized algorithm. Appendix A.6 and A.7 presents details on the condition and examples of Ω -bridge construction.
- Appendix B shows additional experiment details and results.

A THEORETICAL ANALYSIS ON BRIDGES

A.1 DERIVATION OF THE MAIN LOSS IN EQUATION (13)

[Proof of Equation (13)] Denote by $\mathbb{Q}^x = \mathbb{Q}(\cdot | Z_T = x)$. Note that

$$\begin{aligned} \mathcal{KL}(\mathbb{Q}^{\Pi^*} \parallel \mathbb{P}^\theta) &= \mathbb{E}_{x \sim \Pi^*, Z \sim \mathbb{Q}^x} \left[\log \frac{d\mathbb{Q}^{\Pi^*}}{d\mathbb{P}^\theta}(Z) \right] \\ &= \mathbb{E}_{x \sim \Pi^*, Z \sim \mathbb{Q}^x} \left[\log \frac{d\mathbb{Q}^x}{d\mathbb{P}^\theta}(Z) + \log \frac{d\mathbb{Q}^{\Pi^*}}{d\mathbb{Q}^x}(Z) \right] \\ &= \mathbb{E}_{x \sim \Pi^*} [\mathcal{KL}(\mathbb{Q}^x \parallel \mathbb{P}^\theta)] + \text{const}, \end{aligned}$$

where *const* denotes a constant that is independent of θ . Recall that \mathbb{Q}^x follows $dZ_t = \eta^x(Z_{[0,t]}, t)dt + \sigma(Z_t, t)dW_t$, and \mathbb{P}^θ follows $dZ_t = s^\theta(Z_t, t)dt + \sigma(Z_t, t)dW_t$. By Girsanov theorem (e.g., Lejay, 2018),

$$\begin{aligned} \mathcal{KL}(\mathbb{Q}^x \parallel \mathbb{P}^\theta) &= \mathcal{KL}(\mathbb{Q}_0^x \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \mathbb{E}_{Z \sim \mathbb{Q}^x} \left[\int_0^T \|s^\theta(Z_t, t) - \eta^x(Z_{[0,t]}, t)\|_2^2 dt \right] \\ &= \mathbb{E}_{Z \sim \mathbb{Q}^x} \left[-\log p_0^\theta(Z_0) + \frac{1}{2} \int_0^T \|s^\theta(Z_t, t) - \eta^x(Z_{[0,t]}, t)\|_2^2 dt \right] + \text{const}. \end{aligned}$$

Hence

$$\begin{aligned} L(\theta) &= \mathbb{E}_{x \sim \Pi^*, Z \sim \mathbb{Q}^x} \left[-\log p_0^\theta(Z_0) + \frac{1}{2} \int_0^T \|s^\theta(Z_t, t) - \eta^x(Z_{[0,t]}, t)\|_2^2 dt \right] + \text{const} \\ &= \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[-\log p_0^\theta(Z_0) + \frac{1}{2} \int_0^T \|s^\theta(Z_t, t) - \eta^{Z^x}(Z_{[0,t]}, t)\|_2^2 dt \right] + \text{const}. \end{aligned}$$

A.2 DERIVATION OF THE DRIFT η^{Π^*} OF \mathbb{Q}^{Π^*}

Lemma A.1. *Let \mathbb{Q}^x is the law of*

$$dZ_t^x = \eta^x(Z_{[0,t]}^x, t)dt + \sigma(Z_t^x, t)dW_t, \quad Z_0 \sim \mathbb{Q}_0^x,$$

and $\mathbb{Q}^{\Pi^*} := \int \mathbb{Q}^x(Z)\Pi^*(dx)$ for a distribution Π^* on \mathbb{R}^d . Then \mathbb{Q}^{Π^*} is the law of

$$dZ_t = \eta^{\Pi^*}(Z_{[0,t]}, t)dt + \sigma(Z_t, t)dW_t, \quad Z_0 \sim \mathbb{Q}_0^{\Pi^*},$$

where

$$\eta^{\Pi^*}(z_{[0,t]}, t) = \mathbb{E}_{x \sim \Pi^*, Z \sim \mathbb{Q}^x} [\eta^x(Z_{[0,t]}, t) | Z_{[0,t]} = z_{[0,t]}], \quad \mathbb{Q}_0^{\Pi^*}(dz_0) = \mathbb{E}_{x \sim \Pi^*} [\mathbb{Q}_0^x(dz_0)].$$

[Proof] \mathbb{Q}^{Π^*} is the solution of the following optimization problem:

$$\mathbb{Q}^{\Pi^*} = \arg \min_{\mathbb{P}} \left\{ \mathcal{KL}(\mathbb{Q}^{\Pi^*} \parallel \mathbb{P}) = \mathbb{E}_{x \sim \Pi^*} [\mathcal{KL}(\mathbb{Q}^x \parallel \mathbb{P})] + \text{const} \right\}.$$

By Girsanov’s Theorem (e.g., [Lejay, 2018](#)), any stochastic process \mathbb{P} that has $\mathcal{KL}(\mathbb{Q}^x \parallel \mathbb{P}) < +\infty$ (and hence is equivalent to \mathbb{Q}^x) has a form of $dZ_t = \eta^{\Pi^*}(Z_{[0,t]}, t)dt + \sigma(Z_t, t)dW_t$ for some measurable function η^{Π^*} , and

$$\begin{aligned} & \mathbb{E}_{x \sim \Pi^*}[\mathcal{KL}(\mathbb{Q}^x \parallel \mathbb{P})] \\ &= \mathbb{E}_{x \sim \Pi^*}[\mathcal{KL}(\mathbb{Q}_0^x \parallel \mathbb{P}_0)] + \mathbb{E}_{x \sim \Pi^*, Z \sim \mathbb{Q}^x} \left[\frac{1}{2} \int_0^T \left\| \sigma(Z_t, t)^{-1} (\eta^{\Pi^*}(Z_{[0,t]}, t) - \eta^x(Z_{[0,t]}, 0)) \right\|_2^2 \right]. \end{aligned}$$

It is clear that to achieve the minimum, we need to take $\mathbb{P}_0(\cdot) = \mathbb{E}_{x \sim \Pi^*}[\mathbb{Q}_0^x(\cdot)]$ and $\eta^{\Pi^*}(z_{[0,t]}, t) = \mathbb{E}_{x \sim \Pi^*, Z \sim \mathbb{Q}^x}[\eta^x(Z_{[0,t]}, t) \mid Z_{[0,t]} = z_{[0,t]}]$, which yields the desirable form of \mathbb{Q}^{Π^*} .

A.3 \mathbb{P}^{θ^*} YIELDS A MARKOVIZATION OF \mathbb{Q}^{Π^*}

As \mathbb{P}^{θ} is Markov by the model assumption, it can not perfectly fit \mathbb{Q}^{Π^*} which is non-Markov in general. This is a substantial problem because \mathbb{Q}^{Π^*} can be non-Markov *even if* \mathbb{Q}^x is Markov for all $x \in \Omega$ (see Section A.4). In fact, using Doob’s h -transform method ([Doob, 1984](#)), \mathbb{Q}^{Π^*} can be shown to be the law of a diffusion process

$$dZ_t = \eta^{\Pi^*}(Z_{[0,t]}, t)dt + \sigma(Z_t, t)dW_t, \quad \eta^{\Pi^*}(z_{[0,t]}, t) = \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}}[\eta^{Z_T}(z_{[0,t]}, t) \mid Z_{[0,t]} = z_{[0,t]}],$$

where η^{Π^*} is the expectation of η^x when $x = Z_T$ is drawn from \mathbb{Q} conditioned on $Z_{[0,t]}$.

We resolve this by observing that it is not necessary to match the whole path measure ($\mathbb{P}^{\theta} \approx \mathbb{Q}^{\Pi^*}$) to match the terminal ($\mathbb{P}_T^{\theta} \approx \mathbb{Q}_T^{\Pi^*} = \Pi^*$). It is enough for \mathbb{P}^{θ} to be the best Markov approximation (a.k.a. Markovization) of \mathbb{Q}^{Π^*} , which matches all (hence terminal) fixed-time marginals with \mathbb{Q}^{Π^*} :

$$\text{Proj}(\mathbb{Q}^{\Pi^*}, \mathcal{M}) := \arg \min_{\mathbb{P} \in \mathcal{M}} \mathcal{KL}(\mathbb{Q}^{\Pi^*} \parallel \mathbb{P}), \quad \mathcal{M} = \text{the set of all Markov processes on } [0, T].$$

Proposition A.2. *The global optimum of $\mathcal{L}(\theta)$ in (11) and (13) is achieved by θ^* if*

$$s^{\theta^*}(z, t) = \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}}[\eta^{Z_T}(Z_{[0,t]}, t) \mid Z_t = z], \quad \mu^{\theta^*}(dz_0) = \mathbb{Q}_0^{\Pi^*} = \mathbb{E}_{x \sim \Pi^*}[\mathbb{Q}_0^x(dz_0)]. \quad (14)$$

In this case, $\mathbb{P}^{\theta^*} = \text{Proj}(\mathbb{Q}^{\Pi^*}, \mathcal{M})$ is the Markovization of \mathbb{Q}^{Π^*} , with which it matches all time-marginals: $\mathbb{P}_t^{\theta^*} = \mathbb{Q}_t^{\Pi^*}$ for all time $t \in [0, T]$. In addition,

$$\mathcal{KL}(\Pi^* \parallel \mathbb{P}_T^{\theta^*}) \leq \mathcal{KL}(\mathbb{P}^{\theta^*} \parallel \mathbb{P}^{\theta}) = \mathcal{KL}(\mathbb{Q}^{\Pi^*} \parallel \mathbb{P}^{\theta}) - \mathcal{KL}(\mathbb{Q}^{\Pi^*} \parallel \mathbb{P}^{\theta^*}) = \mathcal{L}(\theta) - \mathcal{L}(\theta^*). \quad (15)$$

Note that s^{θ^*} is a conditional expectation of η^{Π^*} : $s^{\theta^*}(z, t) = \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}}[\eta^{\Pi^*}(Z_{[0,t]}, t) \mid Z_t = z]$. Theorem 1 of [Peluchetti \(2021\)](#) gives a related result that the marginals of mixtures of Markov diffusion processes can be matched by another Markov diffusion process, but does not discuss the issue of Markovization nor connect to KL divergence. Theorem 1 of [Song et al. \(2021\)](#) is the special case of (15) when \mathbb{Q}^{Π^*} is Markov.

[Proof of Proposition A.2] It is the combined result of Lemma A.3 and Lemma A.4 below.

Lemma A.3. *Let \mathbb{Q} be a non-Markov diffusion process on $[0, T]$ of form*

$$\mathbb{Q}: \quad dZ_t = \eta(Z_{[0,t]}, t)dt + \sigma(Z_t, t)dW_t, \quad Z_0 \sim \mathbb{Q}_0,$$

and $\mathbb{M} = \arg \min_{\mathbb{P} \in \mathcal{M}} \mathcal{KL}(\mathbb{Q} \parallel \mathbb{P})$ be the Markovization of \mathbb{Q} , where \mathcal{M} is the set of all Markov processes on $[0, T]$. Then \mathbb{Q} is the law of

$$\mathbb{M}: \quad dZ_t = m(Z_t, t)dt + \sigma(Z_t, t)dW_t, \quad Z_0 \sim \mathbb{Q}_0,$$

where

$$m(z, t) = \mathbb{E}_{Z \sim \mathbb{Q}}[\eta(Z_{[0,t]}, t) \mid Z_t = z].$$

In addition, we have $\mathbb{Q}_t = \mathbb{M}_t$ for all time $t \in [0, T]$.

[Proof] By Girsanov's Theorem (e.g., [Lejay, 2018](#)), any process that has $\mathcal{KL}(\mathbb{Q} \parallel \mathbb{M}) < +\infty$ (and hence is equivalent to \mathbb{Q}) has a form of $dZ_t = m(Z_{[0,t]}, t)dt + \sigma(Z_t, t)dW_t$, where m is a measurable function. Since \mathbb{M} is Markov, we have $m(Z_{[0,t]}, t) = m(Z_t, t)$. Then

$$\mathcal{KL}(\mathbb{Q} \parallel \mathbb{P}) = \mathcal{KL}(\mathbb{Q}_0 \parallel \mathbb{P}_0) + \mathbb{E}_{Z \sim \mathbb{Q}} \left[\frac{1}{2} \int_0^T \|\sigma(Z_t, t)^{-1}(\eta(Z_{[0,t]}, t) - m(Z_t, 0))\|_2^2 \right].$$

It is clear that to achieve the minimum, we need to take $\mathbb{M}_0 = \mathbb{Q}_0$ and $m(z, t) = \mathbb{E}_{Z \sim \mathbb{Q}}[\eta(Z_{[0,t]}, t) \mid Z_t = z]$.

To prove $\mathbb{Q}_t = \mathbb{M}_t$, note that by the chain rule of KL divergence:

$$\mathcal{KL}(\mathbb{Q} \parallel \mathbb{P}) = \mathcal{KL}(\mathbb{Q}_t \parallel \mathbb{P}_t) + \mathbb{E}_{Z_t \sim \mathbb{Q}_t} [\mathcal{KL}(\mathbb{Q}(\cdot \mid Z_t) \parallel \mathbb{P}(\cdot \mid Z_t))], \quad \forall t \in [0, T].$$

As the second term $\mathbb{P}(\cdot \mid Z_t)$ is independent of the choice of the marginal \mathbb{P}_t at time $t \in [0, T]$, the optimum should be achieved by \mathbb{M} only if $\mathbb{M}_t = \mathbb{Q}_t$.

Lemma A.4. *Let*

$$\begin{aligned} \mathbb{Q} : \quad & dZ_t = \eta(Z_{[0,t]}, t)dt + \sigma(Z_t, t)dW_t, \quad Z_0 \sim \mathbb{Q}_0 \\ \mathbb{M} : \quad & dZ_t = m(Z_t, t)dt + \sigma(Z_t, t)dW_t, \quad Z_0 \sim \mathbb{Q}_0, \\ \mathbb{P}^\theta : \quad & dZ_t = s^\theta(Z_t, t)dt + \sigma(Z_t, t)dW_t, \quad Z_0 \sim \mathbb{P}_0^\theta, \end{aligned}$$

where \mathbb{M} is the Markovization of \mathbb{Q} (see [Lemma A.3](#)). Then

$$\mathcal{KL}(\mathbb{Q} \parallel \mathbb{P}^\theta) = \mathcal{KL}(\mathbb{Q} \parallel \mathbb{M}) + \mathcal{KL}(\mathbb{M} \parallel \mathbb{P}^\theta).$$

Hence, assume there exists θ^* such that $\mathbb{P}^{\theta^*} = \mathbb{M}$ and write $\mathcal{L}(\theta) := \mathcal{KL}(\mathbb{Q} \parallel \mathbb{P}^\theta)$. We have

$$\mathcal{KL}(\mathbb{Q}_T \parallel \mathbb{P}_T^\theta) = \mathcal{KL}(\mathbb{M}_T \parallel \mathbb{P}_T^\theta) \leq \mathcal{KL}(\mathbb{M} \parallel \mathbb{P}^\theta) = \mathcal{L}(\theta) - \mathcal{L}(\theta^*).$$

[Proof] Note that

$$\begin{aligned} & \mathcal{KL}(\mathbb{M} \parallel \mathbb{P}^\theta) \\ &= \mathcal{KL}(\mathbb{M}_0 \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \mathbb{E}_{Z_t \sim \mathbb{M}_t} \left[\int_0^T \|\sigma(Z_t, t)^{-1}(s^\theta(Z_t, t) - m(Z_t, t))\|_2^2 \right] dt \\ &= \mathcal{KL}(\mathbb{M}_0 \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \int_0^T \mathbb{E}_{Z_t \sim \mathbb{M}_t} \left[\|\sigma(Z_t, t)^{-1}(s^\theta(Z_t, t) - m(Z_t, t))\|_2^2 \right] dt \\ &= \mathcal{KL}(\mathbb{Q}_0 \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \int_0^T \mathbb{E}_{Z_t \sim \mathbb{Q}_t} \left[\|\sigma(Z_t, t)^{-1}(s^\theta(Z_t, t) - m(Z_t, t))\|_2^2 \right] dt \quad // \mathbb{Q}_t = \mathbb{M}_t \quad \forall t \\ &= \mathcal{KL}(\mathbb{Q}_0 \parallel \mathbb{P}_0^\theta) + \mathbb{E}_{Z \sim \mathbb{Q}} \left[\frac{1}{2} \int_0^T \|\sigma(Z_t, t)^{-1}(s^\theta(Z_t, t) - m(Z_t, t))\|_2^2 dt \right] \\ &= \mathcal{KL}(\mathbb{Q}_0 \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \|s^\theta - m\|_{\mathbb{Q}, \sigma}^2, \end{aligned}$$

where we define $\|f\|_{\mathbb{Q}, \sigma}^2 = \mathbb{E}_{Z \sim \mathbb{Q}} \left[\frac{1}{2} \int_0^T \|\sigma(Z_t, t)^{-1}f(Z_t, t)\|_2^2 dt \right]$.

On the other hand,

$$\begin{aligned} \mathcal{KL}(\mathbb{Q} \parallel \mathbb{P}^\theta) &= \mathcal{KL}(\mathbb{Q}_0 \parallel \mathbb{P}_0^\theta) + \mathbb{E}_{Z \sim \mathbb{Q}} \left[\frac{1}{2} \int_0^T \|\sigma(Z_t, t)^{-1}(s^\theta(Z_t, 0)) - \eta(Z_{[0,t]}, t)\|_2^2 dt \right] \\ &= \mathcal{KL}(\mathbb{Q}_0 \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \|s^\theta - \eta\|_{\mathbb{Q}, \sigma}^2 \\ \mathcal{KL}(\mathbb{Q} \parallel \mathbb{M}) &= \mathbb{E}_{Z \sim \mathbb{Q}} \left[\frac{1}{2} \int_0^T \|\sigma(Z_t, t)^{-1}(\eta(Z_{[0,t]}, t) - m(Z_t, 0))\|_2^2 dt \right] \\ &= \frac{1}{2} \|\eta - m\|_{\mathbb{Q}, \sigma}^2. \end{aligned}$$

Using Lemma A.5 with $a(z) = \sigma(z, t)^{-1} s^\theta(z, t)$, and $b(z_{[0,t]}) = \sigma(z, t)^{-1} \eta(z_{[0,t]}, t)$, we have the following bias-variance decomposition:

$$\|\eta - s^\theta\|_{\mathbb{Q}, \sigma}^2 = \|s^\theta - m\|_{\mathbb{Q}, \sigma}^2 + \|\eta - m\|_{\mathbb{Q}, \sigma}^2.$$

Hence, $\mathcal{KL}(\mathbb{Q} \parallel \mathbb{P}^\theta) = \mathcal{KL}(\mathbb{M} \parallel \mathbb{P}^\theta) + \mathcal{KL}(\mathbb{Q} \parallel \mathbb{M})$.

Finally, $\mathcal{KL}(M_T \parallel \mathbb{P}_T^\theta) \leq \mathcal{KL}(\mathbb{M} \parallel \mathbb{P}^\theta)$ is the direct result of the following factorization of KL divergence:

$$\mathcal{KL}(\mathbb{M} \parallel \mathbb{P}^\theta) = \mathcal{KL}(\mathbb{M}_T \parallel \mathbb{P}_T^\theta) + \mathbb{E}_{x \sim \mathbb{M}_T} [\mathcal{KL}(\mathbb{M}_T(\cdot | Z_T = x) \parallel \mathbb{P}_T^\theta(\cdot | Z_T = x))].$$

Lemma A.5. *Let (X, Y) be a random variable and $a(x)$, $b(x, y)$ are square integral functions. Let $m(x) = \mathbb{E}[b(X, Y) | X = x]$. We have*

$$\mathbb{E}[\|a(X) - b(X, Y)\|_2^2] = \mathbb{E}[\|a(X) - m(X)\|_2^2] + \mathbb{E}[\|b(X, Y) - m(X)\|_2^2].$$

[Proof]

$$\begin{aligned} \mathbb{E}[\|a(X) - b(X, Y)\|_2^2] &= \mathbb{E}[\|a(X) - m(X) + m(X) - b(X, Y)\|_2^2] \\ &= \mathbb{E}[\|a(X) - m(X)\|_2^2] + \mathbb{E}[\|m(X) - b(X, Y)\|_2^2] + 2\Delta, \end{aligned}$$

where

$$\begin{aligned} \Delta &= \mathbb{E}[(a(X) - m(X))^\top (m(X) - b(X, Y))] \\ &= \mathbb{E}[(a(X) - m(X))^\top \mathbb{E}[(m(X) - b(X, Y)) | X]] \\ &= \mathbb{E}[(a(X) - m(X))^\top (m(X) - m(X))] = 0. \end{aligned}$$

A.4 MARKOV AND RECIPROCAL PROPERTIES OF \mathbb{Q}^{Π^*}

Mixture of Bridges and Initialization It is an immediate observation that the mixtures of a set of bridges are also bridges: let $\mathbb{Q}^{z, A}$ be a set of A -bridges indexed by a variable z , then $\mathbb{Q}^A := \int \mathbb{Q}^{z, A} \mu(dz)$ is an x -bridge for any distribution μ on z .

A special case is to take the mixture of the conditional bridges in (5) starting from different deterministic initialization, which shows that we can obtain a valid x -bridge by equipping the same drift in (5) with essentially *any* initialization. Hence, the choices of the drift force and initialization in \mathbb{Q}^x can be completely decouple.

Proposition A.6. *Let $\tilde{\mathbb{Q}}$ is a path measure and Ω_x is the set of z for which $\tilde{\mathbb{Q}}^{z_0, x}(\cdot) := \tilde{\mathbb{Q}}(\cdot | Z_T = x, Z_0 = z_0)$ exists. Then $\mathbb{Q}^x := \int \tilde{\mathbb{Q}}^{z_0, x} \mu(dz_0 | x)$ is an x -bridge, for any distribution μ on $\Omega \times \Omega$.*

[Proof of Proposition A.6] This is an obvious result. We have $\mathbb{Q}^{z_0, x}(Z_T = x) = 1$ by the definition of conditioned processes. Hence $\mathbb{Q}^x(Z_T = x) = \int \mathbb{Q}^{z_0, x}(Z_T = x) \mu(dz_0 | x) = \int \mu(dz_0 | x) = 1$.

Markov and Reciprocal Properties of \mathbb{Q}^{Π^*} If \mathbb{Q}^x is constructed as $\mathbb{Q}^x = \mathbb{Q}(\cdot | Z_T = x)$, it is easy to see that $\mathbb{Q}^{\Pi^*} := \int \mathbb{Q}^x(\cdot) \Pi^*(dx)$ is Markov iff \mathbb{Q} is Markov.

Proposition A.7. *Assume $\mathbb{Q}^x = \mathbb{Q}(\cdot | Z_T = x)$ and $\pi^*(z) := \frac{d\Pi^*}{d\mathbb{Q}_T}(z)$ exists and is positive everywhere. Then \mathbb{Q}^{Π^*} is Markov, iff \mathbb{Q} is Markov.*

[Proof of Proposition A.7] If $\mathbb{Q}^x = \mathbb{Q}(\cdot | Z_T = x)$, we have from the definition of \mathbb{Q}^{Π^*} :

$$\mathbb{Q}^{\Pi^*}(Z) = \mathbb{Q}(Z | Z_T) \Pi^*(Z_T) = \mathbb{Q}(Z) \pi^*(Z_T),$$

where $\pi^*(Z_T) = \frac{d\Pi^*}{d\mathbb{Q}^x}(Z_T)$. Therefore, \mathbb{Q}^{Π^*} is obtained by multiplying a positive factor $\pi^*(Z_T)$ on the terminal state Z_T of \mathbb{Q} . Hence \mathbb{Q}^{Π^*} has the same Markov structure as that of \mathbb{Q} .

If \mathbb{Q}^x is constructed from mixtures of bridges as above, the resulting \mathbb{Q}^{Π^*} is more complex. In fact, simply varying the initialization μ in Proposition (A.6) can change the Markov structure of \mathbb{Q}^{Π^*} .

Proposition A.8. *Take \mathbb{Q}^x to be the dynamics in (7) initialized from $Z_0 \sim \mathcal{N}(0, v_0)$. Assume $\sigma_t > 0$, $\forall t \in [0, T]$. Then \mathbb{Q}^{Π^*} is Markov only when $v_0 = 0$, or $v_0 = +\infty$.*

[Proof of Proposition A.8] When taking \mathbb{Q}^x to be the dynamics (7) initialized from $Z_0 \sim \mu_0 = \mathcal{N}(0, v_0)$, we have $\mathbb{Q}^x = \int \mu_0(dz_0) \tilde{\mathbb{Q}}^{z_0, x}$, where $\tilde{\mathbb{Q}}^{z_0, x} = \tilde{\mathbb{Q}}(\cdot | Z_0 = z_0, Z_T = x)$ with $\tilde{\mathbb{Q}}$ following Brownian motion $dZ_t = dW_t$. Hence, we can write $\mathbb{Q}^{\Pi^*}(dZ) = \tilde{\mathbb{Q}}(dZ)r(Z_0, Z_T)$, where $r(z_0, z_T) = \frac{d\mu_0 \otimes \Pi^*}{d\tilde{\mathbb{Q}}_{0, T}}(z_0, z_T)$. From Léonard et al. (2014), \mathbb{Q}^{Π^*} is Markov iff $r(x, z_0) = f(x)g(z_0)$ for some f and g , which is not the case except the degenerated case ($v_0 = 0$ and $v_0 = +\infty$) because $\tilde{\mathbb{Q}}_{0, 1}$ is not factorized.

On the other hand, when $v_0 = 0$, we have that $\mathbb{Q}^x = \mathbb{Q}(\cdot | Z_T = x)$ is the standard Brownian bridge and hence \mathbb{Q}^{Π^*} is Markov following Proposition A.7. When $v_0 = +\infty$, as the case of SMLD, \mathbb{Q}^{Π^*} is the law of $Z_t = \tilde{Z}_{T-t}$ with $d\tilde{Z}_t = dW_t$ and $\tilde{Z}_0 \sim \Pi^*$, which is also Markov.

The right characterization of \mathbb{Q}^{Π^*} from Proposition (A.6) involves reciprocal processes (Léonard et al., 2014).

Definition A.9. *A process Z with law \mathbb{Q} on $[0, T]$ is said to be reciprocal if it can be written into $\mathbb{Q} = \int \tilde{\mathbb{Q}}^{z_0, z_T} \mu(dz_0, dz_T)$, where $\tilde{\mathbb{Q}}$ is a Markov process and $\tilde{\mathbb{Q}}^{z_0, z_T} = \tilde{\mathbb{Q}}(\cdot | Z_0 = z_0, Z_T = z_T)$, and μ is a probability measure on $\Omega \times \Omega$.*

Proposition A.10. *\mathbb{Q}^{Π^*} is reciprocal iff $\mathbb{Q}^x = \int \tilde{\mathbb{Q}}^{z_0, x} \mu(dz_0 | x)$ for a Markov $\tilde{\mathbb{Q}}$ and distribution μ .*

[Proof of Proposition A.10] Note that

$$\mathbb{Q}^{\Pi^*}(\cdot) = \int \pi(dx) \mathbb{Q}^x(\cdot) = \int \pi(dx) \mu(dz_0 | x) \tilde{\mathbb{Q}}^{z_0, x}(\cdot).$$

Hence if $\tilde{\mathbb{Q}}$ is Markov, \mathbb{Q}^{Π^*} is reciprocal by Definition A.9.

On the other hand, if \mathbb{Q}^{Π^*} is reciprocal, we have $\mathbb{Q}^{\Pi^*}(\cdot) = \int \mathbb{M}^{z_0, x}(\cdot) \mu(dz_0, dx)$ for some Markov process \mathbb{M} and probability measure μ on $\Omega \times \Omega$. In this case, we have $\mathbb{Q}^x(\cdot) = \mathbb{Q}^{\Pi^*}(\cdot | Z_T = x) = \int \mathbb{M}^{z_0, x}(\cdot) \mu(dz_0 | x)$, assuming it exists.

Intuitively, a reciprocal process can be viewed as connecting the head and tail of a Markov chain, yielding a single loop structure. A characteristic property is $\mathbb{Q}(X_{[s, t]} \in A | Z_{[0, s]}, Z_{[t, T]}) = \mathbb{Q}(X_{[s, t]} \in A | Z_s, Z_t)$, where A is any event that occur between time s and t . Solutions of the Schrodinger bridge problems are reciprocal processes (Léonard et al., 2014).

A.5 PRACTICAL ALGORITHM AND ERROR ANALYSIS

In practice, we need to introduce empirical and numerical approximations in both training and inference phases. Denote by $\tau = \{\tau_i\}_{i=1}^{K+1}$ a grid of time points with $0 = \tau_1 < \tau_2 \dots < \tau_{K+1} = T$. During training, we minimize an empirical and time-discretized surrogate of $\mathcal{L}(\theta)$ as follows

$$\hat{\mathcal{L}}(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(\theta; Z^{(i)}, \tau^{(i)}), \quad \ell(\theta; Z, \tau) := -\log p_0^\theta(Z_0) + \frac{1}{2K} \sum_{k=1}^K \Delta(\theta; Z, \tau_k), \quad (16)$$

where $\Delta(\theta; Z, t) := \|\sigma^{-1}(Z_t, t)(s^\theta(Z_t, t) - \eta^x(Z_{[0, t]}, t))\|^2$, and $\{Z^{(i)}\}$ is drawn from \mathbb{Q}^{Π^*} , and $\tau^{(i)}$ can be either a deterministic uniform grid of $[0, T]$, i.e., $\tau^{(i)} = \{i/K\}_{i=0}^K$, or drawn i.i.d. uniformly on $[0, T]$ (see e.g., Song et al. (2020b); Ho et al. (2020)). A subtle problem here is that the

variance of $\Delta(\theta; Z, t)$ grows to infinite as $t \uparrow T$. Hence, we should not include $\Delta(\theta; Z, T)$ at the end point $\tau^{K+1} = T$ into the sum in the loss $\ell(\theta, Z, \tau)$ to avoid variance exploding.

In the sampling phase, the continuous-time model \mathbb{P}^θ should be approximated numerically. A standard approach is the Euler-Maruyama method, which simulates the trajectory on a time grid τ by

$$\hat{Z}_{\tau_{k+1}} = \hat{Z}_{\tau_k} + \epsilon_k s^\theta(\hat{Z}_{\tau_k}, \tau_k) + \sqrt{\epsilon_k} \sigma(\hat{Z}_{\tau_k}, \tau_k) \xi_k, \quad \epsilon_k = \tau_{k+1} - \tau_k, \quad \xi_k \sim \mathcal{N}(0, I_d), \quad (17)$$

The final output is \hat{Z}_T . The following result shows the KL divergence between Π^* and the distribution of \hat{Z}_T can be bounded by the sum of the step size and the expected optimality gap $\mathbb{E}[\hat{\mathcal{L}}(\theta) - \hat{\mathcal{L}}(\theta^*)]$ of the time-discretized loss in (16).

A.5.1 TIME-DISCRETIZATION ERROR ANALYSIS (PROPOSITION A.11)

Proposition A.11. Assume $\Omega = \mathbb{R}^d$ and $\sigma(z, t) = \sigma(t)$ is state-independent. Take the uniform time grid $\tau^{\text{unif}} := \{i\epsilon\}_{i=0}^K$ with step size $\epsilon = T/K$ in the sampling step (17). Assume $\sigma(t) > c > 0$, $\forall t$ and $\sigma(t)$ is piecewise constant w.r.t. time grid τ^{unif} . Let $\mathcal{L}_\epsilon(\theta) = \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}}[\ell(\theta; Z, \tau^{\text{unif}})]$. Let $\mathbb{P}_T^{\theta, \epsilon}$ be the distribution of the resulting sample \hat{Z}_T . Let θ^* be an optimal parameter satisfying (14). Assume $C_0 := \sup_{z, t} \left(\|s^{\theta^*}(z, t)\|^2 / (1 + \|z\|^2), \text{tr}(\sigma^2(z, t)), \mathbb{E}_{\mathbb{P}^{\theta^*}}[\|Z_0\|^2] \right) < +\infty$, and $\|s^{\theta^*}(z, t) - s^{\theta^*}(z', t')\|_2^2 \leq L \left(\|z - z'\|^2 + |t - t'| \right)$ for $\forall z, z' \in \mathbb{R}^d$ and $t, t' \in [0, T]$. Then

$$\sqrt{\mathcal{KL}(\Pi^* \parallel \mathbb{P}_T^{\theta, \epsilon})} \leq \sqrt{\mathcal{L}_\epsilon(\theta) - \mathcal{L}_\epsilon(\theta^*)} + O(\sqrt{\epsilon}).$$

We provide the analysis and proof for Proposition A.11 in the following text.

Proposition A.12. Assume $\Omega = \mathbb{R}^d$ and $\sigma(z, t) = \sigma(t)$ is state-independent and $\sigma(t) > c > 0$, $\forall t \in [0, T]$. Take the uniform time grid $\tau^{\text{unif}} := \{i\epsilon\}_{i=0}^K$ with step size $\epsilon = T/K$ in the sampling step (17). Let $\mathcal{L}_\epsilon(\theta) = \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}}[\ell_\epsilon(\theta; Z)]$ with

$$\ell_\epsilon(\theta, Z_t) = -\log p_0^\theta(Z_0) + \frac{1}{2K} \sum_{k=1}^K \left\| \sigma_k^{-1} (s^\theta(Z_{t_k}, t_k) - \eta^{Z_T}(Z_{[0, t_k]}, t_k)) \right\|_2^2,$$

where $\epsilon > 0$ is a step size with $T = K\epsilon$ and $t_k = (k-1)\epsilon$, and $\sigma_k^2 := (t_{k+1} - t_k)^{-1} \int_{t_k}^{t_{k+1}} \sigma(t)^2 dt$. Let $\mathbb{P}_T^{\theta, \epsilon}$ be the distribution of the sample \hat{Z}_T resulting from the following Euler method:

$$\hat{Z}_{t_{k+1}} = \hat{Z}_{t_k} + \epsilon s^\theta(\hat{Z}_{t_k}, t_k) + \sqrt{\epsilon} \sigma_k \xi_k,$$

where $\xi_k \sim \mathcal{N}(0, I_d)$ is the standard Gaussian noise in \mathbb{R}^d . Let θ^* be an optimal parameter satisfying (14). Assume $C_0 := \sup_{z, t} \left(\|s^{\theta^*}(z, t)\|^2 / (1 + \|z\|^2), \text{tr}(\sigma^2(z, t)), \mathbb{E}_{\mathbb{P}^{\theta^*}}[\|Z_0\|^2] \right) < +\infty$, and s^{θ^*} satisfies $\|s^{\theta^*}(z, t) - s^{\theta^*}(z', t')\|_2^2 \leq L \left(\|z - z'\|^2 + |t - t'| \right)$ for $\forall z, z' \in \mathbb{R}^d$ and $t, t' \in [0, T]$. Then we have

$$\sqrt{\mathcal{KL}(\Pi^* \parallel \mathbb{P}_T^{\theta, \epsilon})} \leq \sqrt{\mathcal{L}_\epsilon(\theta) - \mathcal{L}_\epsilon(\theta^*)} + O(\sqrt{\epsilon}).$$

[Proof of Proposition A.11] This is the result of Lemma A.13 below by noting that the $\hat{\mathbb{P}}^\theta$ there is equivalent to the Euler method above, and $\mathcal{L}_\epsilon(\theta) - \mathcal{L}_\epsilon(\theta^*) \leq \tilde{\mathcal{L}}_\epsilon(\theta) - \tilde{\mathcal{L}}_\epsilon(\theta^*)$ (because $\sigma_k^{-2} = ((t_{k+1} - t_k)^{-1} \int_{t_k}^{t_{k+1}} \sigma(t)^2 dt)^{-1} \leq (t_{k+1} - t_k)^{-1} \int_{t_k}^{t_{k+1}} \sigma(t)^{-2} dt$).

Lemma A.13. Let h be a step size and $\epsilon = T/K$ for a positive integer K . For each $t \in [0, \infty)$, denote by $\lfloor t \rfloor_\epsilon = \max(\{k\epsilon : k \in \mathbb{N}\} \cap [0, t])$. Assume

$$\begin{aligned} \mathbb{Q}^{\Pi^*} : \quad & dZ_t = \eta^{\Pi^*}(Z_{[0, t]}, t) dt + \sigma(Z_t, t) dW_t, \quad Z_0 \sim \mathbb{Q}_0 \\ \mathbb{P}^{\theta^*} : \quad & dZ_t = s^{\theta^*}(Z_t, t) dt + \sigma(Z_t, t) dW_t, \quad Z_0 \sim \mathbb{Q}_0, \\ \mathbb{P}^\theta : \quad & dZ_t = s^\theta(Z_t, t) dt + \sigma(Z_t, t) dW_t, \quad Z_0 \sim \mathbb{P}_0^\theta \\ \hat{\mathbb{P}}^\theta : \quad & dZ_t = s^\theta(Z_{\lfloor t \rfloor_\epsilon}, t) dt + \sigma(Z_t, t) dW_t, \quad Z_0 \sim \mathbb{P}_0^\theta, \end{aligned}$$

where \mathbb{P}^{θ^*} is the Markovianization of \mathbb{Q}^{Π^*} , and $\hat{\mathbb{P}}^\theta$ is a discretized version of \mathbb{P}^θ . Define

$$\tilde{\mathcal{L}}_\epsilon(\theta) = \mathbb{E}_{\mathbb{Q}^{\Pi^*}} \left[-\log p_0^\theta(Z_0) + \frac{1}{2} \int_0^T \left\| \sigma^{-1}(Z_t, t) (s^\theta(Z_{[t]_\epsilon}, [t]_\epsilon) - \eta^{Z_T}(Z_{[0, [t]_\epsilon}, [t]_\epsilon)) \right\|^2 dt \right].$$

Assume the conditions of Lemma A.17 holds for \mathbb{P}^{θ^*} , and $\sigma(z, t) \geq c > 0$ for all z, t , and s^θ satisfies $\|s^\theta(z, t) - s^\theta(z', t')\|_2^2 \leq L (\|z - z'\|^2 + |t - t'|)$ for $\forall z, z' \in \mathbb{R}^d$ and $t, t' \in [0, T]$. Then

$$\sqrt{\mathcal{KL}(\mathbb{P}^{\theta^*} \parallel \hat{\mathbb{P}}^\theta)} \leq \sqrt{\tilde{\mathcal{L}}_\epsilon(\theta) - \tilde{\mathcal{L}}_\epsilon(\theta^*)} + C\sqrt{\epsilon},$$

where C is a constant that is independent of ϵ .

[Proof] Define $\|f\|_{\mathbb{Q}, \sigma}^2 = \mathbb{E}_{Z \sim \mathbb{Q}} [\int_0^T \|\sigma(Z, t) f(Z, t)\|^2]$ for convenient notation. Let $s_\epsilon^\theta(Z, t) = s^\theta(Z_{[t]_\epsilon}, [t]_\epsilon)$, and $\eta_\epsilon = \eta^{Z_T}(Z_{[0, [t]_\epsilon}, [t]_\epsilon)$.

$$\begin{aligned} & \mathcal{KL}(\mathbb{P}^{\theta^*} \parallel \hat{\mathbb{P}}^\theta) \\ &= \mathcal{KL}(\mathbb{P}_0^{\theta^*} \parallel \hat{\mathbb{P}}_0^\theta) + \frac{1}{2} \|s^{\theta^*} - s_\epsilon^\theta\|^2 \\ &\leq \mathcal{KL}(\mathbb{P}_0^{\theta^*} \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \left((1 + \omega) \|s_\epsilon^\theta - s_\epsilon^{\theta^*}\|_{\mathbb{P}^{\theta^*}, \sigma}^2 + (1 + 1/\omega) \|s^{\theta^*} - s_\epsilon^{\theta^*}\|_{\mathbb{P}^{\theta^*}, \sigma}^2 \right) \\ &:= (1 + \omega)I_1 + (1 + 1/\omega)I_2, \end{aligned}$$

where $\omega > 0$ is any positive number and

$$\begin{aligned} I_1 &:= \frac{1}{1 + \omega} \mathcal{KL}(\mathbb{P}_0^{\theta^*} \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \|s_\epsilon^\theta - s_\epsilon^{\theta^*}\|_{\mathbb{P}^{\theta^*}, \sigma}^2 \\ &= \frac{1}{1 + \omega} \mathcal{KL}(\mathbb{P}_0^{\theta^*} \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \|s_\epsilon^\theta - s_\epsilon^{\theta^*}\|_{\mathbb{Q}^{\Pi^*}, \sigma}^2 \\ &\leq \mathcal{KL}(\mathbb{P}_0^{\theta^*} \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \|s_\epsilon^\theta - s_\epsilon^{\theta^*}\|_{\mathbb{Q}^{\Pi^*}, \sigma}^2 \\ &\leq \mathcal{KL}(\mathbb{P}_0^{\theta^*} \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \left(\|s_\epsilon^\theta - \eta_\epsilon^{Z_T}\|_{\mathbb{Q}^{\Pi^*}, \sigma}^2 - \|s_\epsilon^{\theta^*} - \eta_\epsilon^{Z_T}\|_{\mathbb{Q}^{\Pi^*}, \sigma}^2 \right) \quad // \text{Lemma A.5} \\ &= \tilde{\mathcal{L}}_\epsilon(\theta) - \tilde{\mathcal{L}}_\epsilon(\theta^*), \end{aligned}$$

and

$$\begin{aligned} I_2 &:= \frac{1}{2} \|s^{\theta^*} - s_\epsilon^{\theta^*}\|_{\mathbb{P}^{\theta^*}, \sigma}^2 \\ &\leq \frac{L}{2} \mathbb{E}_{\mathbb{P}^{\theta^*}} \left[\int_0^T \sigma^{-2}(Z_t, t) (\|Z_t - Z_{[t]_\epsilon}\|^2 + (t - [t]_\epsilon)) dt \right] \\ &\leq \frac{L}{2c^2} \mathbb{E}_{\mathbb{P}^{\theta^*}} \left[\int_0^T ((Z_t - Z_{[t]_\epsilon})^2 + (t - [t]_\epsilon)) dt \right] \\ &\leq \frac{L}{2c^2} (C_{\mathbb{P}^{\theta^*}} + 1) \int_0^T (t - [t]_\epsilon) dt \quad // \text{Lemma A.17} \\ &= \frac{L}{2c^2} (C_{\mathbb{P}^{\theta^*}} + 1) \frac{T\epsilon}{2}. \quad // \text{Lemma A.15,} \end{aligned}$$

where $C_{\mathbb{P}^{\theta^*}}$ is a constant depending on \mathbb{P}^{θ^*} that comes from Lemma A.17. Hence

$$\begin{aligned} \mathcal{KL}(\mathbb{P}^{\theta^*} \parallel \hat{\mathbb{P}}^\theta) &\leq \inf_{\omega \geq 0} (1 + \omega)I_1 + (1 + 1/\omega)I_2 \\ &= (\sqrt{I_1} + \sqrt{I_2})^2 \\ &\leq \left(\sqrt{\tilde{\mathcal{L}}_\epsilon(\theta) - \tilde{\mathcal{L}}_\epsilon(\theta^*)} + \frac{1}{2} \sqrt{\frac{L}{c^2} (C_{\mathbb{P}^{\theta^*}} + 1) T\epsilon} \right). \end{aligned}$$

This completes the proof.

Lemma A.14. For any $a, b \in \mathbb{R}^d$, and $\omega \geq 0$,

$$\|a + b\|_2^2 \leq (1 + \omega) \|a\|_2^2 + (1 + 1/\omega) \|b\|_2^2.$$

[Proof]

$$(1 + \omega) \|a\|_2^2 + (1 + 1/\omega) \|b\|_2^2 \geq \|a\|_2^2 + \|b\|_2^2 + 2a^\top b = \|a + b\|_2^2$$

Lemma A.15. Assume $T \geq 0$, $\epsilon \geq 0$ and $T/\epsilon \in \mathbb{N}$. We have

$$\int_0^T (t - \lfloor t \rfloor_\epsilon) dt = \frac{T\epsilon}{2}.$$

[Proof]

$$\begin{aligned} \int_0^T (t - \lfloor t \rfloor_\epsilon) dt &= \sum_{k=0}^{K-1} \int_0^\epsilon (hk + x - hk) dx \\ &= \sum_{k=0}^{K-1} \int_0^\epsilon x dx \\ &= K\epsilon^2/2 \\ &= T\epsilon/2. \end{aligned}$$

Lemma A.16 (Grönwall's inequality). Let I denote an interval of the real line of the form $[a, \infty)$ or $[a, b]$ or $[a, b)$ with $a < b$. Let α, β and u be real-valued functions defined on I . Assume that β and u are continuous and that the negative part of α is integrable on every closed and bounded subinterval of I .

(a) If β is non-negative and if u satisfies the integral inequality

$$u(t) \leq \alpha(t) + \int_a^t \beta(s)u(s) ds, \quad \forall t \in I,$$

which is true if

$$u'(t) \leq \alpha'(t) + \beta(t)u(t).$$

then

$$u(t) \leq \alpha(t) + \int_a^t \alpha(s)\beta(s) \exp\left(\int_s^t \beta(r) dr\right) ds, \quad t \in I.$$

(b) If, in addition, the function α is non-decreasing, then

$$u(t) \leq \alpha(t) \exp\left(\int_a^t \beta(s) ds\right), \quad t \in I.$$

Lemma A.17. Consider

$$dX_t = b(Z_t, t)dt + \sigma(Z_t, t)dW_t, \quad X_0 = 0, t \in [0, T].$$

Assume there exists a finite constant C_0 , such that

$$\begin{aligned} \|b(x, t)\|_2^2 &\leq C_0(1 + \|x\|_2^2), \quad \forall x \in \mathbb{R}^d, t \in [0, T], \\ \text{tr}(\sigma\sigma^\top(x, t)) &\leq C_0, \quad \forall x \in \mathbb{R}^d, t \in [0, T]. \end{aligned}$$

and $\mathbb{E}[\|Z_0\|_2^2] \leq C_0$.

Then for any $0 \leq s \leq t \leq T$, we have

$$\mathbb{E}[\|Z_t - Z_s\|_2^2] \leq K_{C_0, T}(t - s),$$

where $K_{C_0, T}$ is a finite constant that depends on C_0 and T .

[Proof] Let $\eta = \sup_{x,t} \text{tr}(\sigma\sigma^\top(x,t))$. We have by Ito Lemma,

$$\begin{aligned} \frac{d}{dt} \mathbb{E} \left[\|Z_t - Z_s\|_2^2 \right] &= \mathbb{E} \left[2(Z_t - Z_s)^\top (b(Z_t, t) + dW_t) + \eta \right] \\ &= \mathbb{E} \left[2(Z_t - Z_s)^\top b(Z_t, t) + \eta \right] \\ &= \mathbb{E} \left[\|Z_t - Z_s\|_2^2 + \|b(Z_t, t)\|_2^2 + d \right] \\ &\leq \mathbb{E} \left[\|Z_t - Z_s\|_2^2 + C_0(1 + \|Z_t\|_2^2) + \eta \right] \\ &\leq (1 + 2C_0) \mathbb{E} \left[\|Z_t - Z_s\|_2^2 \right] + \eta + C_0(1 + 2\mathbb{E}[\|Z_s\|_2^2]). \end{aligned}$$

Using Gronwall's inequality,

$$\mathbb{E} \left[\|Z_t - Z_s\|_2^2 \right] \leq (t-s)(\eta + C_0(1 + 2\mathbb{E}[\|Z_s\|_2^2])) \exp((t-s)(1 + 2C_0)).$$

Taking $s = 0$ yields that

$$\mathbb{E} \left[\|Z_t - Z_0\|_2^2 \right] \leq t(\eta + C_0(1 + 2\mathbb{E}[\|Z_0\|_2^2])) \exp(t(1 + 2C_0)).$$

Hence

$$\begin{aligned} \mathbb{E}[\|Z_t\|_2^2] &\leq 2\mathbb{E} \left[\|Z_t - Z_0\|_2^2 \right] + 2\mathbb{E} \left[\|X_0\|_2^2 \right] \\ &\leq 2t(\eta + C_0(1 + 2\mathbb{E}[\|Z_0\|_2^2])) \exp(t(1 + 2C_0)) + 2\mathbb{E} \left[\|Z_0\|_2^2 \right] \\ &\leq 4T(C_0 + C_0^2) \exp(T(1 + 2C_0)) + 2C_0. \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{E} \left[\|Z_t - Z_s\|_2^2 \right] &\leq (t-s)(\eta + C_0(1 + 2\mathbb{E}[\|Z_s\|_2^2])) \exp((t-s)(1 + 2C_0)) \\ &\leq C(t-s), \end{aligned}$$

where

$$C = (2C_0 + 4C_0^2 + 8C_0T(C_0 + C_0^2) \exp(T(1 + 2C_0))) \exp(T(1 + 2C_0)).$$

A.5.2 STATISTICAL ERROR ANALYSIS (PROPOSITION A.18)

To provide a simple analysis of the statistical error, we assume that $\hat{\theta}_n = \arg \min_{\theta} \hat{\mathcal{L}}(\theta)$ is an asymptotically normal M-estimator of θ^* following classical asymptotic statistics (Van der Vaart, 2000), with which we can estimate the rate of the excess risk $\mathcal{L}_\epsilon(\hat{\theta}_n) - \mathcal{L}_\epsilon(\theta^*)$ and hence the KL divergence.

Proposition A.18. *Assume the conditions in Proposition (A.11). Assume $\hat{\theta}_n = \arg \min_{\theta} \hat{\mathcal{L}}_\epsilon(\theta)$ with $\hat{\mathcal{L}}_\epsilon(\theta) = \sum_{i=1}^n \ell(\theta; Z^{(i)}, \tau^{\text{unif}})/n$, $Z^{(i)} \sim \mathbb{Q}^{\Pi^*}$. Take \mathbb{Q}^x to be the standard Brownian bridge $dZ_t^x = \frac{x-Z_t^x}{T-t} dt + dW_t$ with $Z_0 \sim \mathcal{N}(0, v_0)$ and $v_0 > 0$. Assume $\sqrt{n}(\hat{\theta}_n - \theta^*) \xrightarrow{d} \mathcal{N}(0, \Sigma_*)$ as $n \rightarrow +\infty$, where Σ_* is the asymptotic covariance matrix of the M estimator $\hat{\theta}_n$. Assume $\mathcal{L}_\epsilon(\theta)$ is second order continuously differentiable and strongly convex at θ^* . Assume Π^* has a finite covariance and admits a density function π that satisfies $\sup_{t \in [0, T]} \mathbb{E}_{\mathbb{Q}^{\Pi^*}} \left[\|\nabla_{\theta} s^{\theta^*}(Z_t, t)\|^2 (1 + \|\nabla \log \pi(Z_T)\|^2 + \text{tr}(\nabla^2 \log \pi(Z_T))) \right] < +\infty$. We have*

$$\mathbb{E} \left[\sqrt{\mathcal{KL}(\Pi^* \parallel \mathbb{P}_T^{\hat{\theta}_n, \epsilon})} \right] = O \left(\sqrt{\frac{\log(1/\epsilon) + 1}{n}} + \sqrt{\epsilon} \right). \quad (18)$$

The expectation in Eq. (18) is w.r.t. the randomness of $\hat{\theta}_n$. The $\log(1/\epsilon)$ factor shows up as the sum of a harmonic series as the variance of $\Delta(\theta; Z, t)$ grows with $O(1/(T-t))$ when $t \uparrow T$. Taking $\epsilon = 1/n$ yields $\mathcal{KL}(\Pi^* \parallel \mathbb{P}_T^{\hat{\theta}_n, \epsilon}) = O(\log n/n)$. If we want to achieve $\mathcal{KL}(\Pi^* \parallel \mathbb{P}_T^{\hat{\theta}_n, \epsilon}) = O(\eta)$, it is sufficient to take $K = T/\epsilon = O(1/\eta)$ steps and $n = O(\log(1/\eta)/\eta)$ data points.

[Proof of Proposition A.18] Let

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{\epsilon}(\theta) := \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} [\ell(\theta; Z)], \quad \hat{\theta}_n = \arg \min_{\theta} \hat{\mathcal{L}}_{\epsilon}(\theta) := \frac{1}{n} \sum_{i=1}^n \ell(\theta; Z^{(i)}),$$

where $\{Z^{(i)}\}_{i=1}^n$ is drawn i.i.d. from \mathbb{Q}^{Π^*} . We assume that $\hat{\theta}_n$ is an asymptotically normal M-estimator, in which case we have

$$\sqrt{n}(\hat{\theta}_n - \theta^*) \xrightarrow{d} \mathcal{N}(0, \Sigma_*),$$

where

$$\Sigma_* = H_*^{-1} V_* H_*^{-1}, \quad H_* = \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} [\nabla_{\theta\theta}^2 \ell(\theta^*; Z)], \quad V_* = \mathbb{E}[\nabla_{\theta} \ell(\theta^*; Z) \nabla_{\theta} \ell(\theta^*; Z)^{\top}],$$

and

$$n\mathbb{E}[(\mathcal{L}(\hat{\theta}_n) - \mathcal{L}(\theta^*))] \asymp \left[\frac{1}{2} \sqrt{n}(\theta_* - \hat{\theta}_n)^{\top} H_* \sqrt{n}(\theta_* - \hat{\theta}_n) \right] \asymp \frac{1}{2} \text{tr}(H_*^{-1} V_*),$$

where $f \asymp g$ denotes that $f - g = o(1)$. We now need to bound $\text{tr}(H_*^{-1} V_*)$. Combining the results in Lemma A.19 and Lemma A.23, we have when $t_k = (k-1)\epsilon$ and $T = K\epsilon$,

$$\text{tr}(H_*^{-1} V_*) = O\left(1 + \frac{1}{K} \sum_{k=1}^K \frac{1}{T - t_k}\right) = O(1 + \log(1/\epsilon)).$$

Hence,

$$\begin{aligned} \mathbb{E}[\sqrt{\mathcal{KL}(\Pi^* \parallel \mathbb{P}_T^{\hat{\theta}_n, \epsilon})}] &= O\left(\mathbb{E}[\sqrt{\mathcal{L}(\hat{\theta}_n) - \mathcal{L}(\theta^*)}] + \sqrt{\epsilon}\right) \\ &= O\left(\sqrt{\mathbb{E}[\mathcal{L}(\hat{\theta}_n) - \mathcal{L}(\theta^*)]} + \sqrt{\epsilon}\right) \\ &= O\left(\sqrt{\frac{\log(1/\epsilon + 1)}{n}} + \sqrt{\epsilon}\right). \end{aligned}$$

Lemma A.19. Assume the conditions in Proposition A.18. Define

$$I_0 = \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[\left\| \nabla \log p_0^{\theta^*}(Z_0) \right\|^2 \right], \quad I_k = \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[\left\| \nabla_{\theta} s^{\theta^*}(Z_{t_k}, t_k) \right\|^2 \text{tr}(\text{cov}(\eta^{Z^T}(Z_{[0, t_k]}, t_k) \mid Z_{t_k})) \right],$$

for $\forall k = 1, \dots, K$. Then

$$\text{tr}(H_*^{-1} V_*)^{1/2} \leq \frac{1}{\lambda_{\min}(H_*)^{1/2}} \left(I_0^{1/2} + \left(\frac{1}{K} \sum_{k=1}^K I_k \right)^{1/2} \right).$$

[Proof] From Lemma A.21, $\text{tr}(H_*^{-1} V_*) \leq (\lambda_{\min}(H_*))^{-1} \text{tr}(V_*)$. Hence we just need to bound $\text{tr}(V_*)$.

$$\begin{aligned}
\text{tr}(V_*)^{1/2} &= \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[\|\nabla_{\theta} \ell(\theta^*, Z)\|_2^2 \right]^{1/2} \\
&\leq \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[\|\nabla_{\theta} \ell(\theta^*, Z)\|_2^2 \right]^{1/2} + \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[\left\| \nabla_{\theta} s^{\theta^*}(Z_{t_k}, t_k) (s^{\theta^*}(Z_{t_k}, t_k) - \eta^{Z_T}(Z_{[0, t_k]}, t_k)) \right\|_2^2 \right]^{1/2} \\
&\leq \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[\|\nabla_{\theta} \ell(\theta^*, Z)\|_2^2 \right]^{1/2} + \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[\left\| \nabla_{\theta} s^{\theta^*}(Z_{t_k}, t_k) \right\|_2^2 \left\| (s^{\theta^*}(Z_{t_k}, t_k) - \eta^{Z_T}(Z_{[0, t_k]}, t_k)) \right\|_2^2 \right]^{1/2} \\
&= \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[\|\nabla_{\theta} \ell(\theta^*, Z)\|_2^2 \right]^{1/2} + \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[\left\| \nabla_{\theta} s^{\theta^*}(Z_{t_k}, t_k) \right\|_2^2 \text{tr}(\text{cov}(\eta^{Z_T}(Z_{[0, t_k]}, t_k) \mid Z_{t_k})) \right]^{1/2} \\
&= I_0^{1/2} + \frac{1}{K} \sum_{k=1}^K I_k^{1/2} \\
&\leq I_0^{1/2} + \sqrt{\frac{1}{K} \sum_{k=1}^K I_k}.
\end{aligned}$$

Lemma A.20. Assume the results in Lemma A.19 and Lemma A.23 hold. Assume

$$\max_{k \in \{1, \dots, K\}} \mathbb{E}_{Z \sim \Pi^*} \left[\left\| \nabla_{\theta} \tilde{s}^{\theta^*}(Z_{t_k}, t_k) \right\|_2^2 \left(1 + \|\nabla \log \pi^*(Z_T)\|_2^2 + \text{tr}(\nabla^2 \log \pi^*(Z_T)) \right) \right] < +\infty,$$

Then for $k = 1, \dots, K$, we have $I_k = O\left(\frac{1}{T-t_k} + 1\right)$.

[Proof] It is a direction application of (20).

Lemma A.21. Let A and B be two $d \times d$ positive semi-definite matrices. Then $\text{tr}(AB) \leq \lambda_{\max}(A)\text{tr}(B)$.

[Proof] Write A into $A = \sum_{i=1}^d \lambda_i u_i u_i^\top$ where λ_i and u_i is the i -th eigenvalue and eigenvectors of A , respectively. Then

$$\text{tr}(AB) = \text{tr}\left(\sum_{i=1}^d \lambda_i u_i u_i^\top B u_i\right) \leq \lambda_{\max}(A) \text{tr}\left(\sum_{i=1}^d u_i u_i^\top B u_i\right) = \lambda_{\max}(A) \text{tr}(B).$$

Controlling the Conditional Variance of the Regression Problem Assume \mathbb{Q}^x is the standard Brownian bridge:

$$\mathbb{Q}^x: \quad dZ_t^x = \frac{x - Z_t^x}{T-t} dt + dW_t, \quad Z_0 \sim \mathcal{N}(0, v_0). \quad (19)$$

In this case, the (ideal) loss function is

$$\mathcal{L}(\theta) = -\mathbb{E}_{X \sim \Pi^*, Z \sim \mathbb{Q}^x} \left[\log p_0^\theta(Z_0) + \frac{1}{2} \int_0^T \|s^\theta(Z_t, t) - Y_t\|^2 dt \right], \quad \text{where } Y_t = \frac{X - Z_t}{1-t}.$$

The second part of the loss is a least square regression for predicting $Y_t = \eta^X(Z_t, t)$ with $s^\theta(Z_t, t)$. The conditioned variance $\text{cov}(Y_t \mid Z_t)$ is an important factor that influences the error of the regression problem. We now show that $\text{tr}(\text{cov}(Y_t \mid Z_t)) = O(1/T - t)$ which means that it explodes to infinity when $t \uparrow T$.

First, note that $\text{tr}(\text{cov}(Y_t \mid Z_t)) = \frac{1}{(T-t)^2} \text{tr}(\text{cov}(X \mid Z_t))$. Using the estimate in Lemma A.23, we have

$$\text{tr}(\text{cov}(Y_t \mid Z_t)) = O\left(\frac{1}{T-t} + \mathbb{E} \left[\|\nabla_x \log \pi^*(X)\|_2^2 + \text{tr}(\nabla^2 \log \pi^*(X)) \mid Z_t \right]\right). \quad (20)$$

Lemma A.22. For the standard Brownian bridge in (19), we have

$$Z_t^x \sim \mathcal{N}\left(\frac{t}{T}x, \frac{t(T-t)}{T} + \frac{(T-t)^2}{T^2}v_0\right).$$

[Proof] Let $Z_t^{z_0, x}$ be the same process that is initialized from $Z_0^{z_0, x} = z_0$. We have from the textbook result regarding Brownian bridge that we can write $Z_t^{z_0, x} = \frac{tx + (T-t)z_0}{T} + \sqrt{\frac{t(T-t)}{T}}\xi_t$ where ξ_t is some standard Gaussian random variable. The result follows directly as $Z_t^x = Z_t^{Z_0, x}$ with $Z_0 \sim \mathcal{N}(0, v_0)$.

Lemma A.23. Let π^* be the density function Π^* on \mathbb{R}^d whose covariance matrix exists. When $X \sim \Pi^*$ and $Z \sim \mathbb{Q}^X$ from (19) with $v_0 > 0$. Then the density function $\rho_t(x|z)$ of $X|Z_t = z_t$ satisfies

$$\rho_t(x|z) \propto \pi^*(x) \exp\left(-\frac{\|\frac{T}{t}z - x\|_2^2}{2\left(\frac{T(T-t)}{t} + v_0\frac{(T-t)^2}{t^2}\right)}\right). \quad (21)$$

In addition, there exists positive constants $c < +\infty$ and $\tau \in (0, T)$, such that

$$\text{tr}(\text{cov}_{\rho_t}(x|z)) \leq \begin{cases} w_t d + w_t^2 \mathbb{E}_{\rho_t} \left[\|\nabla_x \log \pi^*(x)\|_2^2 + \text{tr}(\nabla^2 \log \pi^*(x)) \mid z \right], & \text{when } \tau \leq t \leq T \\ c, & \text{when } 0 \leq t \leq \tau, \end{cases}$$

where $w_t = \frac{T(T-t)}{t} + v_0\frac{(T-t)^2}{t^2}$. So $\text{tr}(\text{cov}_{\rho_t}(x|z))$ is bounded and decay to zero with rate $O(T-t)$ as $t \uparrow T$.

[Proof] We know that $X \sim \Pi^*$ and $Z_t^X|X \sim \mathcal{N}(t/TX, w_t)$. Hence, (21) is a direct result of Bayes rule. Then Lemma A.25 gives

$$\text{tr}(\text{cov}_{\rho_t}(x|z)) = w_t d + w_t^2 \mathbb{E}_{\rho_t} \left[\|\nabla_x \log \pi^*(x)\|_2^2 + \text{tr}(\nabla^2 \log \pi^*(x)) \mid z \right].$$

On the other hand,

$$\rho_t(x|z) \propto \pi^*(x) \exp\left(-\frac{1}{2w_t} \|x\|^2 + \frac{T}{tw_t} z^\top x\right),$$

When $t \rightarrow 0$, we have $1/w_t \rightarrow 0$ and $T/(tw_t) \rightarrow 0$. Hence, $\rho_t(x|z)$ converges to $\pi^*(x)$ as $t \rightarrow 0$, as a result, $\text{tr}(\text{cov}_{\rho_t}(x|z)) \rightarrow \text{tr}(\text{cov}_{\pi^*}(x)) < +\infty$. Therefore, for any $c > 0$, there exists $t_0 > 0$, such that $\text{tr}(\text{cov}_{\rho_t}(x|z)) \leq \text{tr}(\text{cov}_{\pi^*}(x)) + c$ when $0 \leq t \leq t_0$.

Remark A.24. We need to have $v_0 > 0$ to ensure that $T/(tw_t) \rightarrow 0$ in the proof of Lemma A.23. This is purely a technical reason, for yielding a finite bound of the conditioned variance when t is close to 0. We can establish the same result when $v_0 = 0$ by adding the assumption that $\max_{k \in 1, \dots, K} \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[\|\nabla_{\theta} s^{\theta^*}(Z_{t_k}, t_k)\|_2^2 \text{tr}(\text{cov}_{\Pi_{Z_{t_k}}^*}(Z_T)) \right] < +\infty$, where Π_z^* is the distribution with density $\pi_z^*(x) \propto \pi^*(x) \exp(z^\top x/T)$.

Lemma A.25. Let $p(x) \propto \pi(x) \exp\left(-\alpha \frac{\|x-b\|_2^2}{2}\right)$ be a positive probability density function on \mathbb{R}^d , where $\alpha > 0$, $b \in \mathbb{R}$ and $\log \pi$ is continuously second order differentiable. Then

$$\text{tr}(\text{cov}_p(x)) \leq \mathbb{E}_p[\|x\|_2^2] = \alpha^{-1}d + \alpha^{-2} \left(\mathbb{E}_p[\|\nabla_x \log \pi(x)\|_2^2 + \text{tr}(\nabla^2 \log \pi(x))] \right).$$

[Proof] Let us focus on the case when $b = 0$ first. Stein's identity says that

$$\mathbb{E}_p \left[(\nabla_x \log \pi(x) - \alpha x)^\top \phi(x) + \nabla_x^\top \phi(x) \right] = 0,$$

for a general continuously differentiable function ϕ when the integrals above are finite.

Taking $\phi = x$ yields that

$$\mathbb{E}_p [(\nabla_x \log \pi(x) - \alpha x)^\top x + d] = 0,$$

which gives

$$\mathbb{E}_p [\|x\|_2^2] = \alpha^{-1} (\mathbb{E}_p [\nabla_x \log \pi(x)^\top x] + d).$$

On the other hand, taking $\phi(x) = \nabla_x \log \pi(x)$ yields

$$\mathbb{E}_p [(\nabla_x \log \pi(x) - \alpha x)^\top \nabla_x \log \pi(x) + \text{tr}(\nabla^2 \log \pi(x))] = 0,$$

which gives

$$\mathbb{E}_p [\nabla_x \log \pi(x)^\top x] = \alpha^{-1} \left(\mathbb{E}_p [\|\nabla_x \log \pi(x)\|_2^2 + \text{tr}(\nabla^2 \log \pi(x))] \right).$$

This gives

$$\mathbb{E}_p [\|x\|_2^2] = d\alpha^{-1} + \alpha^{-2} \left(\mathbb{E}_p [\|\nabla_x \log \pi(x)\|_2^2 + \text{tr}(\nabla^2 \log \pi(x))] \right).$$

For $b \neq 0$, define $\tilde{p}(x) \propto \pi(x+b) \exp\left(-\frac{\alpha}{2} \|x\|_2^2\right)$, which is the distribution of $\tilde{x} = x - b$ when $x \sim p$. Then applying the result above to \tilde{p} yields

$$\begin{aligned} \text{tr}(\text{cov}_p(x)) &= \text{tr}(\text{cov}_{\tilde{p}}(x)) \\ &\leq \alpha^{-1} d + \alpha^{-2} \mathbb{E}_{x \sim \tilde{p}} \left[\|\nabla_x \log \pi(x+b)\|_2^2 + \text{tr}(\nabla^2 \log \pi(x+b)) \right] \\ &= \alpha^{-1} d + \alpha^{-2} \mathbb{E}_{x \sim p} \left[\|\nabla_x \log \pi(x)\|_2^2 + \text{tr}(\nabla^2 \log \pi(x)) \right]. \end{aligned}$$

A.6 CONDITION FOR Ω -BRIDGES

We provide the proof for Proposition 2.3.

[Proof of Proposition 2.3] By the formula of KL divergence between two diffusion processes, we have

$$\mathcal{KL}(\mathbb{Q}^\Omega \parallel \mathbb{P}^\theta) = \mathcal{KL}(\mathbb{Q}_0^\Omega \parallel \mathbb{P}_0^\theta) + \frac{1}{2} \mathbb{E}_{Z \sim \mathbb{Q}^\Omega} \left[\int_0^T \|f^\theta(Z_t, t)\|_2^2 dt \right] < +\infty.$$

This means that \mathbb{Q}^Ω and \mathbb{P}^θ are absolutely continuous to each other, and hence have the same support. Therefore, $\mathbb{Q}^\Omega(Z_T \in \Omega) = 1$ implies that $\mathbb{P}^\theta(Z_T \in \Omega) = 1$.

A.7 EXAMPLES OF Ω -BRIDGES

If Ω is a product space, the integration can be factorized into one-dimensional integrals. Specifically, assume $\Omega = I_1 \times \dots \times I_d$, then

$$\eta^\Omega(z, t) = [\eta^{I_i}(z_i, t)]_{i=1}^d,$$

where η^{I_i} is the drift fore of the I_i -bridge, and z_i is the i -th element of $z = [z_i]$. Therefore, it is sufficient to focus on 1D case below.

Consider the bridge process constructed from the Brownian motion in (9). If Ω is a discrete set, say $\Omega = \{e_1, \dots, e_K\}$, we have

$$\begin{aligned} \eta^\Omega(z, t) &= \sigma_t^2 \frac{1}{\sum_{k=1}^K \omega(e_k, z, t)} \sum_{k=1}^K \omega(e_k, z, t) \frac{e_k - z}{\beta_T - \beta_t} \\ &= \sigma_t^2 \nabla_z \log \sum_{k=1}^K \omega(e_k, z, t), \end{aligned}$$

age	workclass	fnlwgt	education	marital-status	occupation	relationship	race	sex	capital-gain	capital-loss	hours-per-week	native-country	income
55	Self-emp-not-inc	229791	Bachelors	Divorced	Prof-specialty	Other-relative	White	Male	2561	77	37	Cuba	>50K
45	Self-emp-not-inc	313479	Some-college	Never-married	Machine-op-inspct	Husband	White	Male	4660	483	38	United-States	>50K
81	Without-pay	106624	9th	Married-civ-spouse	Adm-clerical	Husband	Black	Male	1910	66	24	United-States	<=50K
41	Private	285037	9th	Divorced	Adm-clerical	Husband	White	Male	3754	46	28	Poland	<=50K
22	Private	21419	Bachelors	Widowed	Exec-managerial	Own-child	White	Female	4689	209	34	United-States	<=50K

Figure 5: Generated tabular data from Mixed-Bridge.

where

$$\omega(e_k, z, t) = \exp\left(-\frac{\|z - e_k\|^2}{2(\beta_T - \beta_t)}\right).$$

If $\Omega = [a, b]$, we have

$$\begin{aligned} \eta^\Omega(z, t) &= \sigma_t^2 \frac{1}{\int_a^b \omega(e, z, t)} \int_a^b \omega(e, z, t) \frac{e - z}{\beta_T - \beta_t} de \\ &= \sigma_t^2 \nabla_z \log \int_a^b \omega(e, z, t) de \\ &= \sigma_t^2 \nabla_z \log \left(F\left(\frac{z - a}{\sqrt{\beta_T - \beta_t}}\right) - F\left(\frac{z - b}{\sqrt{\beta_T - \beta_t}}\right) \right), \end{aligned}$$

where F is the standard Gaussian CDF.

B ADDITIONAL MATERIALS OF THE EXPERIMENTS

In our experiments, $T = 1$ and $\epsilon = T/K = 1/K$. Moreover, we take the time grid by randomly sampling from $\{i/K\}_{i=0}^{K-1}$ for the training objective Eq. (13). For evaluation, we calculate the standard evidence lower bound (ELBO) by viewing the resulting time-discretized model as a latent variable model:

$$\mathbb{E}_{X \sim \Pi^*} [-\log \hat{p}_T^\theta(X)] \leq \mathbb{E}_{Z \sim \mathbb{Q}^{\Pi^*}} \left[-\log \frac{\hat{p}_0^\theta(Z_0)}{q_0(Z_0)} - \sum_{k=1}^K \log \frac{\hat{p}_{t_{k+1}|t_k}^\theta(Z_{t_{k+1}}|Z_{t_k})}{q_{t_{k+1}|t_k}(Z_{t_{k+1}}|Z_{t_k})} \right],$$

where $t_k = (k-1)\epsilon$, and \hat{p}^θ is the density function of the time-discretized version of \mathbb{P}^θ , and q is the density function of \mathbb{Q} . We adopt Monte-Carlo sampling to estimate the log-likelihood. As in (Song et al., 2020b), we repeat 5 times in the test set for the estimation. For categorical/integer/grid generation, the likelihood of the last step should take the rounding into account: in practice, we have $\hat{Z}_T = \text{rounding}(\hat{Z}_{t_K} + \epsilon s^\theta(\hat{Z}_{t_K}, t_K) + \sqrt{\epsilon} \sigma(Z_{t_K}, t_K) \xi_K, \Omega)$, where $\text{rounding}(x, \Omega)$ denotes finding the nearest element of x on Ω , and hence the likelihood $\hat{p}_{T|t_K}^\theta$ of the last step should incorporate the rounding operator as a part of the model.

B.1 GENERATING MIXED-TYPE TABULAR DATA

In this experiment, the metrics are measured by the implementation from Synthetic Data Vault (SDV) (Patki et al., 2016). For baseline methods, we adopt their open-sourced official implementation¹. For the machine learning models adopted for evaluation, logistic regression, AdaBoost and MLP, we directly use their default configuration in SDV. For the results in Table 1, we repeat the experiments with 5 different random seeds and report their standard deviation. We provide additional generated samples from Mixed-Bridge in Figure 5.

¹<https://github.com/sdv-dev/CTGAN>




Figure 6: Visualization of the noise schedule of Noise decay A, Noise decay B and Noise decay C.

B.2 GENERATING INTEGER-VALUED POINT CLOUDS

In this experiment, we need to process point cloud data on integer grid. To prepare the data, we firstly sample 2048 points from the ground truth mesh. Then, we normalize all the point clouds to a unit bounding box. After this, we simply project the points onto grid point by rounding the coordinate to integer. The metrics in the main text, MMD, COV and 1-NNA are computed with respect to the post-processed integer-valued training point clouds. For the results in Table 2, we repeat the experiments for 3 times and report the mean of the experiments.

B.3 GENERATING SEMANTIC SEGMENTATION MAPS ON CITYSCAPES

In this experiment, we set (*Noise Decay A*): $\sigma_t^2 = 3 \exp(-3t)$; (*Noise Decay B*): $\sigma_t^2 = 3(1 - t)$; (*Noise Decay C*): $\sigma_t^2 = 3 - 3 \exp(-3(1 - t))$. We visualize the noise schedule in Figure 6. Note that, except for Constant Noise, all the other three processes gradually decrease the magnitude of the noise as $t \rightarrow 1$. For fair comparison, we use the same neural network as in Hoogeboom et al. (2021). The network is optimized with Adam optimizer with a learning rate of 0.0002. The model is trained for 500 epochs. The CityScapes dataset (Cordts et al., 2016) contains photos captured by the cameras on the driving cars. A pixel-wise semantic segmentation map is labeled for each photo. As in (Hoogeboom et al., 2021), we rescale the segmentation maps from cityscapes to 32×64 images using nearest neighbour interpolation. Our training set and test set is exactly the same as that of (Hoogeboom et al., 2021) for fair comparison. For the results in Table 3, we repeat the experiments for 3 times and report the mean of the experiments. We provide more samples in Figure 9.

B.4 DISCRETE CIFAR10 GENERATION

The model is trained using the same training strategy as DDPM (Ho et al., 2020) with the code base provided in Song et al. (2020b). Specifically, the neural network is the same U-Net structure as the implementation in Song et al. (2020b). The optimizer is Adam with a learning rate of 0.0002. According to common practice (Song & Ermon, 2020; Song et al., 2020b), the training is smoothed by exponential moving average (EMA) with a factor of 0.999. We use $K = 1000$ and $dt = 0.001$ for discretizing the SDE. To account for the discretization error, after the final step, we apply rounding to the generated images to get real integer-valued images. We compare the value distribution of the generated images in Figure 8.




Figure 7: Diffusion process of one pixel (a 8-dimensional vector) in CityScapes. As $t \rightarrow 1$, 7 of the dimensions reaches 0, while 1 of the dimensions reaches 1, turning the vector into a one-hot vector.




Figure 8: Final value distribution of the generated images with Bridge-Continuous and Bridge-Integer (before rounding) on CIFAR10. We only show the values in $[125.5, 130.5]$ for visual clarity. Integer-Bridge generates discrete values.

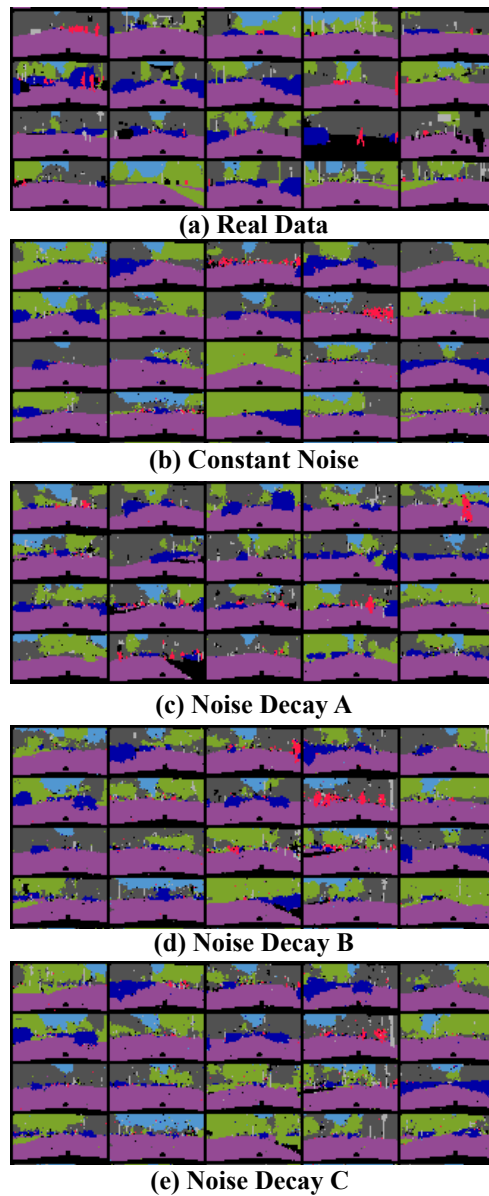


Figure 9: Additional samples from real data, Constant Noise, Noise decay A, Noise decay B and Noise decay C.