

HOW DISCRETE AND CONTINUOUS DIFFUSION MEET: COMPREHENSIVE ANALYSIS OF DISCRETE DIFFUSION MODELS VIA A STOCHASTIC INTEGRAL FRAMEWORK

Yinuo Ren

ICME

Stanford University

yinuoren@stanford.edu

Haoxuan Chen

ICME

Stanford University

haoxuanc@stanford.edu

Grant M. Rotskoff

Department of Chemistry and ICME

Stanford University

rotskoff@stanford.edu

Lexing Ying

Department of Mathematics and ICME

Stanford University

lexing@stanford.edu

ABSTRACT

Discrete diffusion models have gained increasing attention for their ability to model complex distributions with tractable sampling and inference. However, the error analysis for discrete diffusion models remains less well-understood. In this work, we propose a comprehensive framework for the error analysis of discrete diffusion models based on Lévy-type stochastic integrals. By generalizing the Poisson random measure to that with a time-independent and state-dependent intensity, we rigorously establish a stochastic integral formulation of discrete diffusion models and provide the corresponding change of measure theorems that are intriguingly analogous to Itô integrals and Girsanov’s theorem for their continuous counterparts. Our framework unifies and strengthens the current theoretical results on discrete diffusion models and obtains the first error bound for the τ -leaping scheme in KL divergence. With error sources clearly identified, our analysis gives new insight into the mathematical properties of discrete diffusion models and offers guidance for the design of efficient and accurate algorithms for real-world discrete diffusion model applications.

1 INTRODUCTION

Diffusion and flow-based models designed for discrete distributions have gained significant attention in recent years due to their versatility and wide applicability across various domains. These models have been proposed and refined in several key works (Sohl-Dickstein et al., 2015; Austin et al., 2021; Floto et al., 2023; Hooeboom et al., 2021a;b; Meng et al., 2022; Richemond et al., 2022; Sun et al., 2022; Santos et al., 2023). The appeal of such models stems from their potential to address challenging problems in fields like computational biology, where they have shown promise in tasks such as molecule, protein, and DNA sequence design (Seff et al., 2019; Alamdari et al., 2023; Avdeyev et al., 2023; Emami et al., 2023; Frey et al., 2023; Watson et al., 2023; Yang et al., 2023b; Campbell et al., 2024; Stark et al., 2024; Kerby & Moon, 2024; Yi et al., 2024). Additionally, these approaches have proven effective in combinatorial optimization (Li et al., 2024e), modeling retrosynthesis (Igashov et al., 2023), image synthesis (Lezama et al., 2022; Gu et al., 2022), text summarization (Dat et al., 2024) along with the generation of graph (Niu et al., 2020; Shi et al., 2020; Qin et al., 2023; Vignac et al., 2022), layout (Inoue et al., 2023; Zhang et al., 2023), motion (Chi et al., 2024; Lou et al., 2023), sound (Campbell et al., 2022; Yang et al., 2023a), image (Hu et al., 2022; Zhu et al., 2022; Ren et al., 2025a), speech (Wu et al., 2024) and text (He et al., 2022; Wu et al., 2023; Gong et al., 2023; Zheng et al., 2023; Zhou et al., 2023; Shi et al., 2024; Sahoo et al., 2024). Discrete diffusion models also synergize with other methodologies, including tensor networks (Causer et al., 2024), enhanced guidance mechanisms (Gruver et al., 2024; Nisonoff et al., 2024; Li et al., 2024d), structured preferential generation (Rissanen et al., 2024), and alternative metrics, *e.g.* the Fisher

information metric (Davis et al., 2024). These developments highlight the growing importance of discrete modeling in advancing both theoretical understanding and efficient implementations. For a more complete review of related literature, one may refer to the appendix of Ren et al. (2025a).

Partly due to the absence of a discrete equivalent to Girsanov’s theorem, the error analysis for discrete diffusion models remains underdeveloped compared to their continuous counterparts. Existing theoretical work, including a Markov chain-based error analysis for τ -leaping in total variation distance by Campbell et al. (2022) and further advancements for the particular state space $\mathbb{X} = \{0, 1\}^d$ by Chen & Ying (2024), are largely algorithm-specific and not quite easy to be generalized. In this work, our goal is to establish a comprehensive framework for analyzing discrete diffusion models through a stochastic analysis perspective, which is motivated by the theory of continuous diffusion models and different from previous works in many aspects. Drawing on tools from Lévy processes and methodologies for analyzing the simulations of chemical reactions (Li, 2007; Anderson et al., 2011), we extend Poisson random measures to those with evolving intensities, *i.e.* both time-inhomogeneous and state-dependent intensities (Protter, 1983), introduce Lévy-type stochastic integrals (Applebaum, 2009), and articulate corresponding change of measure theorems, which are analogous to the Itô integrals and Girsanov’s theorem in continuous settings.

We further demonstrate that discrete diffusion models implemented via either the τ -leaping or the uniformization scheme can be formulated as stochastic integrals w.r.t. Poisson random measures with evolving intensity, which leads to a unified framework for error analysis. This new stochastic integral-based framework, marking a first for discrete diffusion models, is especially convenient and straightforward for decomposing inference error into three parts: truncation, approximation, and discretization, drawing satisfying parallels with state-of-the-art theories for continuous diffusion model (Chen et al., 2022; 2023a; Benton et al., 2023a). Our approach thus provides intuitive explanations for the loss design and unifies the error analysis across both schemes, enhancing the comparative understanding of their convergence characteristics. Notably, we achieve stronger convergence results in KL divergence and relax some of the stringent assumptions previously required for the state space, the rate matrix, and the estimation of score functions, *etc.*, thereby paving the way for the analysis of a broader class of discrete diffusion models of interest, providing valuable tools for designing efficient and accurate algorithms tailored to the practical demands of discrete diffusion models in real-world applications, and facilitating the transfer of theoretical and practical insights between continuous and discrete diffusion models.

1.1 CONTRIBUTIONS

Our main contributions are summarized as follows:

- We develop a rigorous framework for discrete diffusion models using Lévy-type stochastic integrals based on the Poisson random measure with evolving intensity, which includes formulating discrete diffusion models into stochastic integrals and establishing change of measure theorems that facilitate explicit log-likelihood ratio calculations;
- Our framework extends to a comprehensive, continuous-time analysis for error decomposition in discrete diffusion models, drawing clear parallels with the methodologies used in continuous models and enabling more effective adaptations of techniques across different model types;
- We unify and fortify existing research on discrete diffusion models by deriving the first error bound for τ -leaping in terms of KL divergence, stronger compared to earlier results in TV distance, and providing a comparative study of τ -leaping and uniformization implementations.

1.2 RELATED WORKS

Continuous Diffusion Models. Continuous diffusion models have been one of the most active research areas in generative modeling. Earlier work on continuous diffusion models and probability flow-based models include (Sohl-Dickstein et al., 2015; Zhang et al., 2018; Song & Ermon, 2019; Ho et al., 2020; Song et al., 2020; 2021; Lipman et al., 2022; Liu et al., 2022; Albergo & Vanden-Eijnden, 2022; Albergo et al., 2023). It has shown state-of-the-art performance in various fields of science and engineering. For some recent work and comprehensive review articles, one may refer to (Xu et al., 2022; Yang et al., 2023c; Chan, 2024; Wang et al., 2023; Alakhdar et al., 2024; Chen et al., 2024b; Fan et al., 2024; Guo et al., 2024; Riesel et al., 2024; Zhu et al., 2024).

Theory of Continuous Diffusion Models. In addition to the huge success achieved by diffusion models in empirical studies, many works have also tried to establish sampling guarantees for diffu-

sion and probability flow-based models, such as (Tzen & Raginsky, 2019; Block et al., 2020; Benton et al., 2023b; Chen et al., 2023b; Mbacke & Rivasplata, 2023; Liang et al., 2024; Benton et al., 2024; Ren et al., 2025b). Regarding theoretical analysis of continuous diffusion models, (Lee et al., 2022) provided the first sampling guarantee under the smoothness and isoperimetry assumptions. Follow-up work removed such assumptions (Chen et al., 2022; 2023a; Lee et al., 2023) and obtained better convergence results (Benton et al., 2023a; Pedrotti et al., 2023; Li et al., 2023; 2024a;b). For the probability flow-based implementation, sampling guarantee was also established and further refined in many recent work (Chen et al., 2024c; Gao & Zhu, 2024; Huang et al., 2024; Li et al., 2024c)

2 PRELIMINARIES

In this section, we introduce the basic concepts of both continuous and discrete diffusion models and then roughly outline the error analysis for continuous diffusion models, which will serve as a reference for the error analysis for discrete diffusion models.

2.1 CONTINUOUS DIFFUSION MODELS

In diffusion models, the *forward process* is designed as an Itô process $(\mathbf{x}_t)_{0 \leq t \leq T}$ in \mathbb{R}^d satisfying the following stochastic differential equation (SDE):

$$d\mathbf{x}_t = \mathbf{b}_t(\mathbf{x}_t)dt + \mathbf{g}_t d\mathbf{w}_t, \text{ with } \mathbf{x}_0 \sim p_0, \quad (2.1)$$

where $(\mathbf{w}_t)_{t \geq 0}$ is a standard Brownian motion. The probability distribution of \mathbf{x}_t is denoted by p_t , and the distribution p_0 at time $t = 0$ is the target distribution for sampling. The time-reversal $(\tilde{\mathbf{x}}_s)_{0 \leq s \leq T}$ of (2.1) satisfies the *backward process*:

$$d\tilde{\mathbf{x}}_s = \left[-\tilde{\mathbf{b}}_s(\tilde{\mathbf{x}}_s) + \tilde{\mathbf{g}}_s \tilde{\mathbf{g}}_s^\top \nabla \log \tilde{p}_s(\tilde{\mathbf{x}}_s) \right] ds + \tilde{\mathbf{g}}_s d\mathbf{w}_s, \quad (2.2)$$

where $\tilde{\mathbf{x}}_s$ denotes \mathbf{x}_{T-s} , with $\tilde{p}_0 = p_T$ and $\tilde{p}_T = p_0$.

One of the common choices for the drift \mathbf{b}_t and the diffusion coefficient \mathbf{g} is $\mathbf{b}_t(\mathbf{x}) = -\frac{1}{2}\beta_t \mathbf{x}_t$ and $\mathbf{g} = \sigma\sqrt{\beta_t}\mathbf{I}$, under which (2.1) is an Ornstein-Uhlenbeck (OU) process converging exponentially, i.e. $p_T \approx p_\infty := \mathcal{N}(0, \sigma^2 \mathbf{I})$, and the forward process (2.1) and the backward process (2.2) reduce to the following form:

$$d\mathbf{x}_t = -\frac{1}{2}\beta_t \mathbf{x}_t dt + \sigma\sqrt{\beta_t} d\mathbf{w}_t, \text{ and } d\tilde{\mathbf{x}}_s = \tilde{\beta}_s \left[\frac{1}{2}\tilde{\mathbf{x}}_s + \sigma^2 \nabla \log \tilde{p}_s(\tilde{\mathbf{x}}_s) \right] ds + \sigma\sqrt{\tilde{\beta}_s} d\mathbf{w}_s. \quad (2.3)$$

In practice, the score function $\mathbf{s}_t(\mathbf{x}_t) := \nabla \log p_t(\mathbf{x}_t)$ is often estimated by a neural network $\hat{\mathbf{s}}_t^\theta(\mathbf{x}_t)$, where θ denotes the parameters, and trained via *denoising score-matching* (Hyvärinen & Dayan, 2005; Vincent, 2011):

$$\begin{aligned} \theta &= \arg \min_{\theta} \int_0^T \psi_t \mathbb{E}_{\mathbf{x}_t \sim p_t} \left[\left\| \nabla \log p_t(\mathbf{x}_t) - \hat{\mathbf{s}}_t^\theta(\mathbf{x}_t) \right\|^2 \right] dt \\ &= \arg \min_{\theta} \int_0^T \psi_t \mathbb{E}_{\mathbf{x}_0 \sim p_0} \left[\mathbb{E}_{\mathbf{x}_t \sim p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)} \left[\left\| \nabla \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0) - \hat{\mathbf{s}}_t^\theta(\mathbf{x}_t) \right\|^2 \right] \right] dt, \end{aligned} \quad (2.4)$$

where $p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)$ is the transition distribution from \mathbf{x}_0 to \mathbf{x}_t under (2.3) with an explicit form as

$$\mathcal{N}(\boldsymbol{\mu}_t, \sigma_t^2 \mathbf{I}), \text{ where } \boldsymbol{\mu}_t = \mathbf{x}_0 e^{-\frac{1}{2} \int_0^t \beta_t dt} \text{ and } \sigma_t^2 = \sigma^2 \left(1 - e^{-\int_0^t \beta_t dt} \right), \quad (2.5)$$

and ψ_t is a weighting function for the loss at time t . After obtaining the NN-based score function $\hat{\mathbf{s}}_t^\theta(\mathbf{x}_s)$, the backward process in (2.3) is approximated as:

$$d\mathbf{y}_s = \left[\frac{1}{2}\mathbf{y}_s + \hat{\mathbf{s}}_s^\theta(\mathbf{y}_s) \right] ds + d\mathbf{w}_s, \text{ with } \mathbf{y}_0 \sim q_0 = \mathcal{N}(0, \sigma^2 \mathbf{I}). \quad (2.6)$$

2.2 DISCRETE DIFFUSION MODELS

In discrete diffusion models, one turns to consider a continuous-time Markov chain $(\mathbf{x}_t)_{0 \leq t \leq T}$ in a space \mathbb{X} of finite cardinality as the *forward process*. We denote the probability distribution of \mathbf{x}_t by a vector $\mathbf{p}_t \in \Delta^{|\mathbb{X}|}$, where $\Delta^{|\mathbb{X}|}$ denotes the probability simplex in $\mathbb{R}^{|\mathbb{X}|}$. Given the target distribution \mathbf{p}_0 , the Markov chain satisfies the following master equation:

$$\frac{d\mathbf{p}_t}{dt} = \mathbf{Q}_t \mathbf{p}_t, \text{ where } \mathbf{Q}_t = (Q_t(y, x))_{x, y \in \mathbb{X}} \in \mathbb{R}^{|\mathbb{X}| \times |\mathbb{X}|} \quad (2.7)$$

is the rate matrix at time t . The rate matrix Q_t satisfies the following two conditions:

$$(i) Q_t(x, x) = - \sum_{y \neq x} Q_t(y, x), \forall x \in \mathbb{X}; (ii) Q_t(x, y) \geq 0, \forall x \neq y \in \mathbb{X}.$$

In the following, we will use a shorthand notation \tilde{Q}_t to denote the matrix Q_t with the diagonal elements set to zero. It can be shown that the corresponding backward process is of the same form but with a different rate matrix (Kelly, 2011):

$$\frac{d\tilde{p}_s}{ds} = \tilde{Q}_s \tilde{p}_s, \quad \text{where} \quad \tilde{Q}_s(y, x) = \begin{cases} \frac{\tilde{p}_s(y)}{\tilde{p}_s(x)} \tilde{Q}_s(x, y), & \forall x \neq y \in \mathbb{X}, \\ - \sum_{y' \neq x} \tilde{Q}_s(y', x), & \forall x = y \in \mathbb{X}. \end{cases} \quad (2.8)$$

The rate matrix Q_t is often chosen to possess certain sparse structures such that the forward process converges to a simple distribution that is easy to sample from. Several popular choices include the uniform and absorbing transitions (Lou et al., 2024).

The common practice is to define the score function (or rather the score vector) as $s_t(x) = (s_t(x, y))_{y \in \mathbb{X}} := \frac{p_t}{p_t(x)}$, for any $x \in \mathbb{X}$ and estimate it by a neural network $\hat{s}_t^\theta(x)$, where the neural network θ is trained by minimizing the score entropy (Benton et al., 2022; Lou et al., 2024):

$$\begin{aligned} \theta &= \arg \min_{\theta} \int_0^T \psi_t \mathbb{E}_{x_t \sim p_t} \left[\sum_{y \neq x} \left(-\log \frac{\hat{s}_t^\theta(x, y)}{s_t(x, y)} - 1 + \frac{\hat{s}_t^\theta(x, y)}{s_t(x, y)} \right) s_t(x, y) Q_t(x, y) \right] dt \\ &= \arg \min_{\theta} \int_0^T \psi_t \mathbb{E}_{x_0 \sim p_0} \left[\sum_{y \neq x} \left(-\frac{p_t(y|x_0)}{p_t(x|x_0)} \log \hat{s}_t^\theta(x, y) + \hat{s}_t^\theta(x, y) \right) Q_t(x, y) \right] dt. \end{aligned} \quad (2.9)$$

Similar to the continuous case, the backward process is approximated by the continuous-time Markov chain with the following master equation with $q_0 = p_\infty$ and rate matrix \hat{Q}_s^θ :

$$\frac{dq_s}{ds} = \hat{Q}_s^\theta q_s, \quad \text{where} \quad \hat{Q}_s^\theta(y, x) = \tilde{s}_s^\theta(x, y) \tilde{Q}_s(x, y), \quad \forall x \neq y \in \mathbb{X}. \quad (2.10)$$

and sampling is accomplished by first sampling from the distribution p_∞ and then evolving the Markov chain accordingly.

2.3 ERROR ANALYSIS OF CONTINUOUS DIFFUSION MODELS

Before we proceed to the error analysis of discrete diffusion models, we first review that the error analysis of continuous diffusion models, which is often conducted by considering the following three error terms:

- **Truncation Error:** The error caused by approximating p_T by p_∞ , which is often of the order $\mathcal{O}(d \exp(-T))$ due to exponential ergodicity;
- **Approximation Error:** The error caused by approximating the score function $\nabla \log p_t(x_t)$ by a neural network $\hat{s}_t^\theta(x_t)$, which is often assumed to be of order $\mathcal{O}(\epsilon)$, where ϵ is a small threshold, given a thorough training process;
- **Discretization Error:** The error caused by numerically solving the SDE (2.6) with Euler-Maruyama scheme or other schemes, e.g. exponential integrator (Zhang & Chen, 2022).

The total error is obtained from these three error terms with proper choices of the order of the time horizon T and the design of the numerical scheme. We extract the following theorem from the state-to-the-art theoretical result (Benton et al., 2023a) for later comparison:

Theorem 2.1 (Error Analysis of Continuous Diffusion Models). *Suppose the time discretization scheme $(s_i)_{i \in [0, N]}$ with $s_0 = 0$ and $s_N = T - \delta$ satisfies $s_{k+1} - s_k \leq \kappa(T - s_{k+1})$ for $k \in [0 : N - 1]$. Assume $\text{cov}(p_0) = \mathbf{I}$, and the score function $\nabla \log p_t(x_t)$ is estimated by the neural network $\hat{s}_t^\theta(x_t)$ with ϵ -accuracy, i.e.*

$$\sum_{k=0}^{N-1} (s_{k+1} - s_k) \mathbb{E}_{\tilde{x}_{s_k} \sim \tilde{p}_{s_k}} \left[\left\| \nabla \log \tilde{p}_{s_k}(\tilde{x}_{s_k}) - \tilde{s}_{s_k}^\theta(\tilde{x}_{s_k}) \right\|^2 \right] \leq \epsilon.$$

Then under the following choice of the order of parameters

$$T = \mathcal{O}(\log(d\epsilon^{-1})), \quad \kappa = \mathcal{O}(d^{-1}\epsilon \log^{-1}(d\epsilon^{-1})), \quad N = \mathcal{O}(d\epsilon^{-1} \log^2(d\epsilon^{-1})),$$

we have $D_{\text{KL}}(p_\delta \| \hat{q}_{s_N}) \leq \epsilon$, where \hat{q}_{s_N} is the distribution of the approximate backward process (2.6) implemented with exponential integrator after N steps.

3 STOCHASTIC INTEGRAL FORMULATION OF DISCRETE DIFFUSION MODELS

In this section, we introduce the stochastic integral formulation of discrete diffusion models. The goal is to establish a path evolution equation analogous to Itô integral (or equivalently, stochastic differential equations) with the master equation (2.7) and (2.8) analogous to the Fokker-Planck equation in the continuous case.

3.1 POISSON RANDOM MEASURE WITH EVOLVING INTENSITY

In the following, the Poisson distribution with expectation λ is denoted by $\mathcal{P}(\lambda)$.

Definition 3.1 (Poisson Random Measure with Evolving Intensity). *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and $(\mathbb{X}, \mathcal{B}, \nu)$ be a measure space and $\lambda_t(y)$ is a non-negative predictable process on $\mathbb{R}^+ \times \mathbb{X} \times \Omega$ satisfying for any $T > 0$, $\int_0^T \int_{\mathbb{X}} 1 \vee |y| \vee |y|^2 \lambda_t(y) \nu(dy) dt < \infty$, a.s.. The random measure $N[\lambda](dt, dy)$ on $\mathbb{R}^+ \times \mathbb{X}$ is called a Poisson random measure with evolving intensity $\lambda_t(y)$ if*

- (i) *For any $B \in \mathcal{B}$ and $0 \leq s < t$, $N[\lambda]((s, t] \times B) \sim \mathcal{P}\left(\int_s^t \int_B \lambda_\tau(y) \nu(dy) d\tau\right)$;*
- (ii) *For any $t \geq 0$ and disjoint sets $\{B_i\}_{i \in [n]} \subset \mathcal{B}$, $\{N_t[\lambda](B_i) := N[\lambda]((0, t] \times B_i)\}_{i \in [n]}$ are independent stochastic processes.*

Well-definedness of the Poisson random measure with evolving intensity is non-trivial and we refer readers to Appendix A for more details and discussions. Intuitively, Poisson random measures randomly assign jumps to intervals along the evolution, with the number of points in each region following a Poisson distribution, while those with evolving intensity not only assign jumps but also locations controlled by the intensity function $\lambda_t(y)$. In the following, we will denote the filtration generated by the Poisson random measure $N[\lambda](dt, dy)$ by $(\mathcal{F}_t)_{t \geq 0}$.

The Poisson random measure defined above admits similar properties as the standard Poisson random measure and the Brownian motion. In particular, one can extend the Itô integral to the Lévy-type stochastic integral w.r.t. Poisson random measure with evolving intensity for predictable processes. It also admits Itô isometry, Itô's formula (Theorem A.10), and Lévy's characterization theorem (Theorem A.9), for which we refer readers to Appendix A.2 for details.

Now we turn to the setting of discrete diffusion models, where the state space \mathbb{X} is finite endowed with the natural σ -algebra $\mathcal{B} = 2^{\mathbb{X}}$ and the counting measure $\nu = \sum_{y \in \mathbb{X}} \delta_y$.

Proposition 3.2 (Stochastic Integral Formulation of Discrete Diffusion Models). *The forward process in discrete diffusion models (2.7) can be represented by the following stochastic integral:*

$$x_t = x_0 + \int_0^t \int_{\mathbb{X}} (y - x_{\tau-}) N[\lambda](d\tau, dy), \text{ with } \lambda_\tau(y) = \tilde{Q}_\tau(y, x_{\tau-}), \quad (3.1)$$

and the backward process in discrete diffusion models (2.8) can be represented by the following stochastic integral:

$$\tilde{x}_t = \tilde{x}_0 + \int_0^t \int_{\mathbb{X}} (y - \tilde{x}_{\tau-}) N[\mu](d\tau, dy), \text{ with } \mu_\tau(y) = \tilde{s}_\tau(\tilde{x}_{\tau-}, y) \tilde{Q}_\tau(\tilde{x}_{\tau-}, y), \quad (3.2)$$

where X_{t-} denotes the left limit of a càdlàg process X_t at time t .

The proof of Proposition 3.2 is provided in Appendix A.4. Figure 1 illustrates the stochastic integral in (3.1), comparing two different intensity functions. It is evident that jumps occur more frequently




Figure 1: Example trajectories of stochastic integrals (3.1) w.r.t. Poisson random measure with different evolving intensities. The intensity is chosen as $\lambda_t(y) = 50f_t$ if $|y - x_{t-}| = 1$ or otherwise 0, as shown in dashed lines. Intuitively, λ_t controls the rate of jumps at time t and location y .

in regions where the intensity is higher and less frequently where it is lower. Additionally, the intensity function, which encapsulates key information about the stochastic process, is often much easier and more straightforward to analyze compared to the complex trajectories of the process.

We would like to remark that the stochastic integral formulation in Proposition 3.2 is tantalizingly close to the Itô integral in continuous diffusion models in the form of stochastic differential equations (cf. (2.1) and (2.2)). Recalling that in the continuous case, Girsanov’s theorem is applied to SDEs for deriving the score-matching loss (2.4) and also for the error analysis by associating the loss with the KL divergence, one may wonder if similar techniques can be applied to discrete diffusion models. The following section gives an affirmative answer to this question by providing a change of measure for Poisson random measures with evolving intensity, which is the theoretical foundation of our stochastic integral-based error analysis framework for discrete diffusion models.

3.2 CHANGE OF MEASURE

The following theorem provides a change-of-measure argument for stochastic integrals w.r.t. Poisson random measures with evolving intensity, analogous to Girsanov’s theorem for Itô integrals w.r.t. Brownian motions.

Theorem 3.3 (Change of Measure for Poisson Random Measure with Evolving Density). *Let $N[\lambda](dt, dy)$ be a Poisson random measure with evolving intensity $\lambda_t(y)$ in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and $h_t(y)$ be a positive predictable process on $\mathbb{R}^+ \times \mathbb{X} \times \Omega$. Suppose the following exponential process is a local \mathcal{F}_t -martingale:*

$$Z_t[h] := \exp \left(\int_0^t \int_{\mathbb{X}} \log h_\tau(y) N[\lambda](d\tau \times dy) - \int_0^t \int_{\mathbb{X}} (h_\tau(y) - 1) \lambda_\tau(y) \nu(dy) d\tau \right), \quad (3.3)$$

and \mathbb{Q} is another probability measure on (Ω, \mathcal{F}) such that $\mathbb{Q} \ll \mathbb{P}$ with Radon-Nikodym derivative $d\mathbb{Q}/d\mathbb{P}|_{\mathcal{F}_t} = Z_t[h]$. Then the Poisson random measure $N[\lambda](dt, dy)$ under the measure \mathbb{Q} is a Poisson random measure with evolving intensity $\lambda_t(y)h_t(y)$.

Intuitively, Theorem 3.3 depicts the relations between the Random-Nikodym derivative ($Z_t[h]$) of the probability measures on which two Poisson random measures with different intensities are defined. $Z_t[h]$ also roughly translates to the likelihood ratio between two paths generated with different intensities, with which one could work on stronger convergence results, e.g. KL divergence.

It is straightforward to derive the following corollary, which was derived in (Benton et al., 2022) with a different technique with Feller processes and adopted in (Lou et al., 2024; Chen & Ying, 2024) in the design of loss functions for the neural network training:

Corollary 3.4 (Equivalence between KL Divergence and Score Entropy-based Loss Function). *Let $\tilde{p}_{0:T}$ and $q_{0:T}$ be the path measures of the backward process (2.8) and the approximate backward process (2.10), then it holds that*

$$\begin{aligned} D_{\text{KL}}(\tilde{p}_T \| q_T) &\leq D_{\text{KL}}(\tilde{p}_{0:T} \| q_{0:T}) \\ &= D_{\text{KL}}(\tilde{p}_0 \| q_0) + \mathbb{E} \left[\int_0^T \int_{\mathbb{X}} K \left(\frac{\tilde{s}_\tau^q(\tilde{x}_{\tau-}, y)}{\tilde{s}_\tau(\tilde{x}_{\tau-}, y)} \right) \tilde{s}_\tau(\tilde{x}_{\tau-}, y) \tilde{Q}_\tau(\tilde{x}_{\tau-}, y) \nu(dy) d\tau \right], \end{aligned} \quad (3.4)$$

where $K(x) = x - 1 - \log x \geq 0$, and the expectation is taken w.r.t. paths generated by the backward process (3.2). Consequently, minimizing the loss function (2.9) for discrete diffusion models is equivalent to minimizing the KL divergence between the path measures of the ground truth and the approximate backward process.

Proofs of the change of measure-related arguments above will be provided in Appendix A.3. One should recall that in the continuous case with Itô integrals, the proximity of two paths in KL divergence only requires a small difference between the drift terms by Girsanov’s theorem, and thus the score function can be trained with the mean squared error loss (2.4) (Song et al., 2020), while in the discrete case, the proximity of two paths requires the likelihood ratio to be close to one, accounting for a more complicated score entropy design (2.9) (Benton et al., 2022; Lou et al., 2024).

4 ERROR ANALYSIS OF DISCRETE DIFFUSION MODELS

In this section, we firstly review two different implementations of the discrete diffusion models, namely τ -leaping (Gillespie, 2001) and uniformization (Van Dijk, 1992), derive their stochastic integral formulations as in Proposition 3.2, and provide our main results for their error analysis.

4.1 ALGORITHMS

4.1.1 τ -LEAPING.

A straightforward algorithm for simulating the backward process is to discretize the integral in (3.2) with an Euler-Maruyama scheme. This leads to the τ -leaping algorithm summarized in Algorithm 1. The main idea is to employ a predetermined time discretization scheme and approximate the inference process within each interval using the intensity observed at the time and location corresponding to the start of the interval.

Algorithm 1: τ -Leaping Algorithm for Discrete Diffusion Model Inference

Input: $\hat{y}_0 \sim q_0$, time discretization scheme $(s_i)_{i \in [0:N]}$ with $s_0 = 0$ and $s_N = T - \delta$, intensity function $\hat{\mu}_s^\theta$ defined in Proposition 4.1, and neural network-based score function estimation \hat{s}_t^θ .
Output: A sample $\hat{y}_{s_N} \sim \hat{q}_{t_N}$.
1 **for** $n = 0$ **to** $N - 1$ **do**
2 $\hat{y}_{s_{n+1}} \leftarrow \hat{y}_{s_n} + \sum_{y \in \mathbb{X}} (y - \hat{y}_{s_n}) \mathcal{P}(\hat{\mu}_{s_n}^\theta(y)(s_{n+1} - s_n));$
3 **end**

(4.1)

As shown in the following proposition, τ -leaping can be formulated as a stochastic integral.

Proposition 4.1 (Stochastic Integral Formulation of τ -Leaping). *The τ -leaping algorithm (Algorithm 1) is equivalent to solving the following stochastic integral equation:*

$$\hat{y}_s = \hat{y}_0 + \int_0^s \int_{\mathbb{X}} (y - \hat{y}_{[\tau]-}) N[\hat{\mu}_{[\cdot]}^\theta](d\tau, dy), \quad (4.2)$$

where the evolving intensity $\hat{\mu}_\tau^\theta(y)$ is given by $\hat{\mu}_{[\tau]}^\theta(y) = \tilde{s}_{[\tau]}^\theta(\hat{y}_{[\tau]-}, y) \tilde{Q}_{[\tau]}(\hat{y}_{[\tau]-}, y) = \hat{\mu}_{s_n}^\theta(y)$, in which we used the symbol $[\tau] = s_n$ for $\tau \in [s_n, s_{n+1})$. We will call the process \hat{y}_s the interpolating process of the τ -leaping algorithm and denote the distribution of \hat{y}_s by \hat{q}_s .

In general, the stochastic integral associated with τ -leaping (4.2) is computed using a piecewise constant intensity rather than the original continuous intensity. This approximation introduces discretization error as a trade-off for more efficient implementation. We refer to Remark A.13 for further justification of the τ -leaping algorithm.

4.1.2 UNIFORMIZATION

Another algorithm considered for simulating the backward process in discrete diffusion models is uniformization. The algorithm is summarized in Algorithm 2, in which $\sigma_{(m)}$ denotes the m -th order statistic of the M uniform random variables on $[0, 1]$, and the randomness in (4.3) should be understood as sampling a categorical distribution and updating the state accordingly.

Algorithm 2: Uniformization Algorithm for Discrete Diffusion Model Inference

Input: $\hat{y}_0 \sim q_0$, time discretization scheme $(s_b)_{b \in [0:N]}$ with $s_0 = 0$ and $s_B = T - \delta$, intensity upper bound process $\bar{\lambda}_s$, intensity function $\hat{\mu}_s^\theta$ defined in Proposition 4.1, and neural network-based score function estimation \hat{s}_t^θ .
Output: A sample $x_{s_B} \sim q_{t_B}$.
1 **for** $b = 0$ **to** $B - 1$ **do**
2 $M \sim \mathcal{P}(\bar{\lambda}_{s_{b+1}}(s_{b+1} - s_b)), \sigma_m \sim \text{Unif}([0, 1])$ for $m \in [M];$
3 **for** $m = 1$ **to** M **do**
4 $\hat{y}_{s_b + \sigma_{(m)}} \leftarrow \begin{cases} y, & \text{with prob. } \hat{\mu}_{s_b + \sigma_{(m)}}^\theta(y) / \bar{\lambda}_{s_{b+1}}, \text{ for } y \in \mathbb{X}, \\ \hat{y}_{s_b}, & \text{with prob. } 1 - \sum_{y \in \mathbb{X}} \hat{\mu}_{s_b + \sigma_{(m)}}^\theta(y) / \bar{\lambda}_{s_{b+1}}; \end{cases}$
5 **end**
6 **end**

(4.3)

The main idea is to simulate the backward process by a Poisson random measure with a piecewise constant intensity upper bound process and then sample the behavior of each jump according to the

intensity $\hat{\mu}_s^\theta(y)$ at time s . The uniformization algorithm also admits a stochastic integral formulation, as shown in the following proposition.

Proposition 4.2 (Stochastic Integral Formulation of Uniformization). *Under the block discretization scheme $(s_b)_{b \in [0, B]}$ with $s_0 = 0$ and $s_B = T - \delta$, and for any $s \in (s_b, s_{b+1}]$, we define the piecewise constant intensity upper bound process by $\bar{\lambda}_s = \sup_{s \in (s_b, s_{b+1}]} \int_{\mathbb{X}} \hat{\mu}_s^\theta(y) \nu(dy)$. Then the uniformization algorithm (Algorithm 2) is equivalent to solving the following stochastic integral equation in the augmented measure space $(\mathbb{X} \times [0, \bar{\lambda}], \mathcal{B} \otimes \mathcal{B}([0, \bar{\lambda}]), \nu \otimes m)$:*

$$y_s = y_0 + \int_0^s \int_{\mathbb{X}} \int_{\mathbb{R}} (y - y_{\tau-}) \mathbf{1}_{0 \leq \xi \leq \int_{\mathbb{X}} \hat{\mu}_\tau^\theta(y) \nu(dy)} N[\hat{\mu}^\theta](d\tau, dy, d\xi), \quad (4.4)$$

where the evolving intensity $\hat{\mu}_\tau^\theta(y)$ is given by $\hat{\mu}_\tau^\theta(y) = \tilde{s}_\tau^\theta(\hat{y}_{\tau-}, y) \tilde{Q}_\tau(\hat{y}_{\tau-}, y)$.

Based on Proposition 4.2, one can show that the uniformization algorithm simulates the backward process in discrete diffusion models accurately (cf. Theorem A.15), and the proofs of the claims above will be provided in Appendix A.4.

4.2 ASSUMPTIONS

We need the following assumptions to ensure the well-definedness of discrete diffusion models. For simplicity, we assume the rate matrix \mathbf{Q}_t is time-homogeneous and symmetric, i.e. $\mathbf{Q}_t = \mathbf{Q}$ for any $t \geq 0$. In fact, the results can be easily extended to the time-inhomogeneous case of the family $\mathbf{Q}_t = \beta_t \mathbf{Q}$ with a rescaling factor β_t , and asymmetric cases will be left for future works.

Assumption 4.3 (Regularity of the Rate Matrix). *The rate matrix \mathbf{Q} satisfies:*

- (i) *For any $x, y \in \mathbb{X}$, $Q(x, y) \leq C$ and $\underline{D} \leq -Q(x, x) \leq \overline{D}$ for some constants $C, \underline{D}, \overline{D} > 0$;*
- (ii) *The modified log-Sobolev constant $\rho(\mathbf{Q})$ of the rate matrix \mathbf{Q} (cf. Definition B.6) is lower bounded by $\rho > 0$.*

Statement (i) assumes the regularity of the rate matrix, which is often trivially satisfied in many applications, while Statement (ii) ensures the exponential convergence of the forward process in discrete diffusion models. In general, $\rho(\mathbf{Q})$ may depend on the connectivity and other structures of the corresponding graph $\mathcal{G}(\mathbf{Q})$ (cf. Definition B.1). Such lower bound has been obtained for specific graphs (e.g. Example B.10 and B.11), and general results are in active research (Saloff-Coste, 1997; Bobkov & Tetali, 2006). We refer readers to Appendix B for further discussions on the literature of the modified log-Sobolev constant, as well as its relation to the spectral gap, the mixing time, etc..

Assumption 4.4 (Bounded Score). *The true score function satisfies $s_t(x, y) \lesssim 1 \vee t^{-1}$, while the learned score function satisfies $\hat{s}_s^\theta(x, y) \in (0, M]$, for any $x, y \in \mathbb{X}$.*

The first part on the asymptotic behavior of the true score corresponds to the estimation $\mathbb{E}[\|s_t\|^2] \sim \mathbb{E}[\|x_t - \mu_t\|^2 / \sigma_t^2] \sim 1 \vee t^{-1}$ in the continuous case (Chen et al., 2023a, Assumption 1) and further justification is provided in Remark B.3. The bound on the estimated score can be easily satisfied by adding truncation in post-processing in the implementation of the NN-based score estimator.

Assumption 4.5 (Continuity of Score Function). *For any $t > 0$ and $y \in \mathbb{X}$ such that $Q(x_{t-}, y) > 0$, we have $\left| \frac{\mu_{t+}(y)}{\mu_t(y)} \right| := \left| \frac{p_t(x_{t-})Q(x_{t-}, y)}{p_t(x_t)Q(x_t, y)} - 1 \right| \lesssim 1 \vee t^{-\gamma}$, for some exponent $\gamma \in [0, 1]$.*

Assumption 4.5 corresponds to the Lipschitz continuity of the score function (cf. (Chen et al., 2022, Assumption 1), (Chen et al., 2023a, Assumption 3)) for continuous diffusion models, and is in light of the postulation that adjacent vertices should have close score function and intensity values. In the worse case, assume $Q(x, y) = \Theta(1)$, then a naïve bound would be $\left| \frac{\mu_{t+}(y)}{\mu_t(y)} \right| \lesssim |s_t(x_t, x_{t-})| \lesssim 1 \vee t^{-1}$ with $\gamma = 1$. However, when the initial distribution is both upper and lower bounded, γ may be as small as 0, and we plan to investigate how this (local) continuity of the score function affects the overall performance of discrete diffusion models.

Assumption 4.6 (ϵ -accurate Score Estimation). *The score function $s_t(x_t)$ is estimated by the neural network $\hat{s}_t^\theta(x_t)$ with ϵ -accuracy, i.e.*

$$\sum_{n=0}^{N-1} (s_{n+1} - s_n) \mathbb{E} \left[\int_{\mathbb{X}} K \left(\frac{\tilde{s}_{s_n}^\theta(\tilde{x}_{s_n}, y)}{\tilde{s}_{s_n}(\tilde{x}_{s_n}, y)} \right) \tilde{s}_{s_n}(\tilde{x}_{s_n}, y) \tilde{Q}(\tilde{x}_{s_n}, y) \nu(dy) \right] \leq \epsilon.$$

This assumption assumes the expressive power and sufficient training of the NN-based score estimator and is standard in diffusion model-related theories (Chen et al., 2022; 2023a; Benton et al., 2023a; Chen et al., 2024c).

4.3 ERROR ANALYSIS

Our main results are presented below for each algorithm introduced in Section 4.1.

4.3.1 τ -LEAPING

Theorem 4.7 (Error Analysis of τ -Leaping). *Suppose the time discretization scheme $(s_i)_{i \in [0, N]}$ with $s_0 = 0$ and $s_N = T - \delta$ satisfies for $k \in [0 : N - 1]$, $s_{k+1} - s_k \leq \kappa (1 \vee (T - s_{k+1})^{1+\gamma-\eta})$, where the exponent η satisfies $\gamma < \eta \lesssim 1 - T^{-1}$ when $\gamma < 1$, and $\eta = 1$ when $\gamma = 1$. With Assumptions 4.3, 4.4, 4.5, and 4.6 in place, we have the following error bound with probability $1 - \mathcal{O}(\epsilon)$*

$$D_{\text{KL}}(p_\delta \| \hat{q}_{T-\delta}) \lesssim \exp(-\rho T) \log |\mathbb{X}| + \epsilon + \bar{D}^2 \kappa T,$$

and under the following choice of the order of parameters:

$$T = \mathcal{O}\left(\frac{\log(\epsilon^{-1} \log |\mathbb{X}|)}{\rho}\right), \quad \kappa = \mathcal{O}\left(\frac{\epsilon \rho}{\bar{D}^2 \log(\epsilon^{-1} \log |\mathbb{X}|)}\right), \quad \delta = \begin{cases} 0, & \gamma < 1, \\ \Omega(e^{-\sqrt{T}}), & \gamma = 1, \end{cases} \quad (4.5)$$

we have $D_{\text{KL}}(p_\delta \| \hat{q}_{T-\delta}) \lesssim \epsilon$ with $N = \kappa^{-1}T = \mathcal{O}\left(\frac{\bar{D}^2 \rho^2 \log^2(\epsilon^{-1} \log |\mathbb{X}|)}{\epsilon}\right)$ total steps.

The conclusions and detailed proof (as provided in Appendix C), whose sketch is given in Appendix C.1, are analogous to the error bound for continuous diffusion models (cf. Theorem 2.1), as a summation of the truncation, approximation, and discretization errors as outlined in Section 2.3. We would like to point out the main differences between continuous and discrete diffusion models:

- **Truncation Error:** While the Ornstein-Uhlenbeck process converges exponentially fast in the continuous diffusion models, the exponential convergence of the forward process in discrete diffusion models is non-trivial for general graphs $\mathcal{G}(\mathbf{Q})$, for which the lower bound on the modified log-Sobolev constant ρ is one of the sufficient conditions. In practice, the exponential convergence of the forward process should be verified for the specific problem at hand;
- **Discretization Error:** In continuous diffusion models, the analysis of the discretization error is based on the Itô integral and Girsanov’s theorem, while for discrete diffusion models, the stochastic integral framework, including the Poisson random measure with evolving intensity (cf. Definition 3.1) and change of measure (cf. Theorem 3.3) that we developed above, is employed instead.

Remark 4.8 (Remark on Early Stopping). *As in Assumption 4.4, the true score function may exhibit singular behavior as $s \rightarrow T$, due to possible vacancy in the target distribution \mathbf{p}_0 . To handle this singularity, two different regimes are considered for the time discretization scheme depending on the continuity parameter γ of the score function (Assumption 4.5). The main intuition is that (a) in the worse case $\gamma = 1$, early stopping at time $s = T - \delta$ is necessary; (b) if the target distribution \mathbf{p}_0 is such well-posed (e.g. both upper and lower bounded) and the rate matrix \mathbf{Q} is constructed in a way that the score exhibits certain (local) continuity reflected by $\gamma < 1$, one may choose an appropriate shrinkage η , with which finite discretization error can be achieved with finite steps.*

In Theorem 4.7, the coefficient \bar{D} roughly translates to the dimension d when the discrete diffusion model is applied to $\mathbb{X} = [S]^d$, where S is the number of states along each dimension. For example, when each state is connected to those at a Manhattan distance of 1 with unit weight, we have $C = 1$, $\bar{D} = 2d$, $\log |\mathbb{X}| = d \log S$ hold in Assumption 4.3. Plugging $\bar{D} = \log |\mathbb{X}| = \mathcal{O}(d)$ into the results, we obtain that the total number of steps $N = \tilde{\mathcal{O}}(d^2)$, where $\tilde{\mathcal{O}}$ denotes the order up to logarithmic factors. This recovers the dependency described in (Campbell et al., 2022, Theorem 1) for τ -leaping with a completely different set of techniques, and importantly, our results do not rely on strong assumptions such as a uniform bound on the true score. We also reduce assumption stringency by relating our assumption on the estimation error (Assumption 4.6) more closely to the practical training loss rather than requiring an L^∞ -accurate score estimation error. Most notably, we provide the first convergence guarantees for τ -leaping in KL divergence, which is stronger than the total variation distance, for discrete diffusion models.

4.3.2 UNIFORMIZATION

The error analysis of the uniformization algorithm requires the following modified assumption on the accuracy of the learned score function $\hat{s}_t^\theta(x_t)$:

Assumption 4.6' (ϵ -accurate Learned Score). *The score function $s_t(x_t)$ is estimated by the neural network $\hat{s}_t^\theta(x_t)$ with ϵ -accuracy, i.e.*

$$\mathbb{E} \left[\int_0^{T-\delta} \int_{\mathbb{X}} K \left(\frac{\tilde{s}_s^\theta(\tilde{x}_{s-}, y)}{\tilde{s}_s(\tilde{x}_{s-}, y)} \right) \tilde{s}_s(\tilde{x}_{s-}, y) \tilde{Q}(\tilde{x}_{s-}, y) \nu(dy) \right] ds \leq \epsilon.$$

Theorem 4.9 (Error Analysis of Uniformization). *Suppose the block discretization scheme $(s_b)_{b \in [0, N]}$ with $s_0 = 0$ and $s_N = T - \delta$ satisfies for $k \in [0 : N - 1]$, $s_{k+1} - s_k \leq \kappa (1 \vee (T - s_{k+1}))$. Under Assumptions 4.3, 4.4, 4.5, and 4.6', we have the following error bound*

$$D_{\text{KL}}(p_\delta \| q_{T-\delta}) \lesssim \exp(-\rho T) \log |\mathbb{X}| + \epsilon.$$

Then with $T = \mathcal{O} \left(\frac{\log(\epsilon^{-1} \log |\mathbb{X}|)}{\rho} \right)$ and the early stopping scheme $\delta = \Omega(e^{-T})$, we have

$$D_{\text{KL}}(p_\delta \| q_{T-\delta}) \lesssim \epsilon \text{ with } \mathbb{E}[N] = \mathcal{O} \left(\frac{\overline{D} \log(\epsilon^{-1} \log |\mathbb{X}|)}{\rho} \right) \text{ steps.}$$

The proof of Theorem 4.9 is deferred to Appendix C, where a corresponding sketch is provided in Appendix C.1. Following a similar argument for Theorem 4.7, the dimensionality dependency of the uniformization scheme is $\tilde{\mathcal{O}}(d)$, confirming the result for the special case $\mathbb{X} = \{0, 1\}^d$ in (Chen & Ying, 2024) and $\mathbb{X} = [S]^d$ in (Zhang et al., 2024). Theorem 4.7 and 4.9 offer a direct comparison of the efficiency of the τ -leaping and uniformization implementations for discrete diffusion models. Our proof reveals that the less favorable quadratic dependency in the τ -leaping scheme arises from the truncation error, which is not present in the uniformization scheme, illustrating a possible advantage of the latter in reducing computational complexity.

Recalling the current state-of-the-art result for continuous diffusion models (Theorem 2.1) is $\tilde{\mathcal{O}}(d)$, we conjecture that $\tilde{\mathcal{O}}(d)$ is also the optimal rate in the discrete case. In the continuous case, the linear dependency does not depend on the accurate simulation of the approximate backward process (2.10) (or (3.2)) and is achievable with Euler-Maruyama schemes. The proof was via an intricate stochastic localization argument (Benton et al., 2023a), with which the bound on $\mathbb{E}[\|\nabla^2 \log p_t(x_t)\|^2]$ is improved by a $\mathcal{O}(d)$ -factor from $\mathcal{O}(d^2)$. The corresponding argument for the τ -leaping scheme of discrete diffusion models would be a possible refinement on Proposition C.2, which we believe is of independent interest and will be explored in future work.

5 CONCLUSION

In this paper, we have developed a comprehensive framework for the error analysis of discrete diffusion models. We rigorously introduced the Poisson random measure with evolving intensity and established the Lévy-type stochastic integral alongside change of measure arguments. These advancements not only hold mathematical significance but also facilitate a clear-cut analysis of discrete diffusion models. Moreover, we demonstrated that the inference process can be formulated as a stochastic integral using the Poisson random measure with evolving intensity, allowing the error to be systematically decomposed and optimized by algorithmic design, mirroring the theoretical framework for continuous diffusion models. Our framework unifies the error analysis of discrete diffusion models and provides the first error bounds for the τ -leaping scheme in KL divergence.

Our results lay a theoretical groundwork for the design and analysis of discrete diffusion models, adaptable to broader contexts, such as time-inhomogeneous and non-symmetric rate matrices. Future research directions include exploring the continuum limit of discrete diffusion models in state spaces and time (Winkler et al., 2024; Zhao et al., 2024) and how to accelerate the implementation via parallel sampling (Chung et al., 2023; Shih et al., 2024; Tang et al., 2024; Cao et al., 2024; Selvam et al., 2024; Chen et al., 2024a; Gupta et al., 2024). We hope our work will inspire further research on both the theoretical analysis and practical applications of discrete diffusion models in various fields.

ACKNOWLEDGMENTS

Lexing Ying acknowledges the support of the National Science Foundation under Award No. DMS-2208163.

REFERENCES

- Amira Alakhdar, Barnabas Poczos, and Newell Washburn. Diffusion models in de novo drug design. *Journal of Chemical Information and Modeling*, 2024.
- Sarah Alamdari, Nitya Thakkar, Rianne van den Berg, Alex X Lu, Nicolo Fusi, Ava P Amini, and Kevin K Yang. Protein generation with evolutionary diffusion: sequence is all you need. *BioRxiv*, pp. 2023–09, 2023.
- Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.
- Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.
- David F Anderson, Arnab Ganguly, and Thomas G Kurtz. Error analysis of tau-leap simulation methods. *The Annals of Applied Probability*, 21(6):2226–2262, 2011.
- David Applebaum. *Lévy processes and stochastic calculus*. Cambridge university press, 2009.
- 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.
- Pavel Avdeyev, Chenlai Shi, Yuhao Tan, Kseniia Dudnyk, and Jian Zhou. Dirichlet diffusion score model for biological sequence generation. In *International Conference on Machine Learning*, pp. 1276–1301. PMLR, 2023.
- Joe Benton, Yuyang Shi, Valentin De Bortoli, George Deligiannidis, and Arnaud Doucet. From denoising diffusions to denoising markov models. *arXiv preprint arXiv:2211.03595*, 2022.
- Joe Benton, Valentin De Bortoli, Arnaud Doucet, and George Deligiannidis. Linear convergence bounds for diffusion models via stochastic localization. *arXiv preprint arXiv:2308.03686*, 2023a.
- Joe Benton, George Deligiannidis, and Arnaud Doucet. Error bounds for flow matching methods. *arXiv preprint arXiv:2305.16860*, 2023b.
- Joe Benton, Yuyang Shi, Valentin De Bortoli, George Deligiannidis, and Arnaud Doucet. From denoising diffusions to denoising markov models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(2):286–301, 2024.
- Adam Block, Youssef Mroueh, and Alexander Rakhlin. Generative modeling with denoising auto-encoders and langevin sampling. *arXiv preprint arXiv:2002.00107*, 2020.
- Sergey G Bobkov and Michel Ledoux. On modified logarithmic sobolev inequalities for bernoulli and poisson measures. *Journal of functional analysis*, 156(2):347–365, 1998.
- Sergey G Bobkov and Prasad Tetali. Modified logarithmic sobolev inequalities in discrete settings. *Journal of Theoretical Probability*, 19:289–336, 2006.
- Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. *Advances in Neural Information Processing Systems*, 35:28266–28279, 2022.
- Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, and Tommi Jaakkola. Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design. *arXiv preprint arXiv:2402.04997*, 2024.

- Jiezhong Cao, Yue Shi, Kai Zhang, Yulun Zhang, Radu Timofte, and Luc Van Gool. Deep equilibrium diffusion restoration with parallel sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2824–2834, 2024.
- Luke Causer, Grant M Rotskoff, and Juan P Garrahan. Discrete generative diffusion models without stochastic differential equations: a tensor network approach. *arXiv preprint arXiv:2407.11133*, 2024.
- Stanley H Chan. Tutorial on diffusion models for imaging and vision. *arXiv preprint arXiv:2403.18103*, 2024.
- Haoxuan Chen, Yinuo Ren, Lexing Ying, and Grant M Rotskoff. Accelerating diffusion models with parallel sampling: Inference at sub-linear time complexity. *arXiv preprint arXiv:2405.15986*, 2024a.
- Hongrui Chen and Lexing Ying. Convergence analysis of discrete diffusion model: Exact implementation through uniformization. *arXiv preprint arXiv:2402.08095*, 2024.
- Hongrui Chen, Holden Lee, and Jianfeng Lu. Improved analysis of score-based generative modeling: User-friendly bounds under minimal smoothness assumptions. In *International Conference on Machine Learning*, pp. 4735–4763. PMLR, 2023a.
- Minshuo Chen, Song Mei, Jianqing Fan, and Mengdi Wang. An overview of diffusion models: Applications, guided generation, statistical rates and optimization. *arXiv preprint arXiv:2404.07771*, 2024b.
- Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru R Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. *arXiv preprint arXiv:2209.11215*, 2022.
- Sitan Chen, Giannis Daras, and Alex Dimakis. Restoration-degradation beyond linear diffusions: A non-asymptotic analysis for ddim-type samplers. In *International Conference on Machine Learning*, pp. 4462–4484. PMLR, 2023b.
- Sitan Chen, Sinho Chewi, Holden Lee, Yuanzhi Li, Jianfeng Lu, and Adil Salim. The probability flow ode is provably fast. *Advances in Neural Information Processing Systems*, 36, 2024c.
- Seunggeun Chi, Hyung-gun Chi, Hengbo Ma, Nakul Agarwal, Faizan Siddiqui, Karthik Ramani, and Kwonjoon Lee. M2d2m: Multi-motion generation from text with discrete diffusion models. *arXiv preprint arXiv:2407.14502*, 2024.
- Hyungjin Chung, Jeongsol Kim, Sehui Kim, and Jong Chul Ye. Parallel diffusion models of operator and image for blind inverse problems. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6059–6069, 2023.
- David R Cox. Some statistical methods connected with series of events. *Journal of the Royal Statistical Society: Series B (Methodological)*, 17(2):129–157, 1955.
- Do Huu Dat, Do Duc Anh, Anh Tuan Luu, and Wray Buntine. Discrete diffusion language model for long text summarization. *arXiv preprint arXiv:2407.10998*, 2024.
- Oscar Davis, Samuel Kessler, Mircea Petrache, Avishek Joey Bose, et al. Fisher flow matching for generative modeling over discrete data. *arXiv preprint arXiv:2405.14664*, 2024.
- Persi Diaconis and Laurent Saloff-Coste. Logarithmic sobolev inequalities for finite markov chains. *The Annals of Applied Probability*, 6(3):695–750, 1996.
- Andreas Eberle. Markov processes. *Lecture Notes at University of Bonn*, 2009.
- Andreas Eberle. Stochastic analysis. <http://www.mi.unikoeln.de/stochana/ws1617/Eberle/StochasticAnalysis2015.pdf>, 2015.
- Patrick Emami, Aidan Perreault, Jeffrey Law, David Biagioni, and Peter St John. Plug & play directed evolution of proteins with gradient-based discrete mcmc. *Machine Learning: Science and Technology*, 4(2):025014, 2023.

- Jiahao Fan, Ziyao Li, Eric Alcaide, Guolin Ke, Huaqing Huang, and E Weinan. Accurate conformation sampling via protein structural diffusion. *bioRxiv*, pp. 2024–05, 2024.
- Griffin Floto, Thorsteinn Jonsson, Mihai Nica, Scott Sanner, and Eric Zhengyu Zhu. Diffusion on the probability simplex. *arXiv preprint arXiv:2309.02530*, 2023.
- Nathan C Frey, Daniel Berenberg, Karina Zadorozhny, Joseph Kleinhenz, Julien Lafrance-Vanasse, Isidro Hotzel, Yan Wu, Stephen Ra, Richard Bonneau, Kyunghyun Cho, et al. Protein discovery with discrete walk-jump sampling. *arXiv preprint arXiv:2306.12360*, 2023.
- Xuefeng Gao and Lingjiong Zhu. Convergence analysis for general probability flow odes of diffusion models in wasserstein distances. *arXiv preprint arXiv:2401.17958*, 2024.
- Daniel T Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of chemical physics*, 115(4):1716–1733, 2001.
- Paul Glasserman and Nicolas Merener. Convergence of a discretization scheme for jump-diffusion processes with state-dependent intensities. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 460(2041):111–127, 2004.
- Sharad Goel. Modified logarithmic sobolev inequalities for some models of random walk. *Stochastic processes and their applications*, 114(1):51–79, 2004.
- Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and Lingpeng Kong. Diffuseq-v2: Bridging discrete and continuous text spaces for accelerated seq2seq diffusion models. *arXiv preprint arXiv:2310.05793*, 2023.
- Leonard Gross. Logarithmic sobolev inequalities. *American Journal of Mathematics*, 97(4):1061–1083, 1975.
- Nate Gruver, Samuel Stanton, Nathan Frey, Tim GJ Rudner, Isidro Hotzel, Julien Lafrance-Vanasse, Arvind Rajpal, Kyunghyun Cho, and Andrew G Wilson. Protein design with guided discrete diffusion. *Advances in neural information processing systems*, 36, 2024.
- Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pp. 10696–10706, 2022.
- Zhiye Guo, Jian Liu, Yanli Wang, Mengrui Chen, Duolin Wang, Dong Xu, and Jianlin Cheng. Diffusion models in bioinformatics and computational biology. *Nature reviews bioengineering*, 2(2):136–154, 2024.
- Shivam Gupta, Linda Cai, and Sitan Chen. Faster diffusion-based sampling with randomized mid-points: Sequential and parallel. *arXiv preprint arXiv:2406.00924*, 2024.
- Zhengfu He, Tianxiang Sun, Kuanning Wang, Xuanjing Huang, and Xipeng Qiu. Diffusionbert: Improving generative masked language models with diffusion models. *arXiv preprint arXiv:2211.15029*, 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Emiel Hoogetboom, Alexey A Gritsenko, Jasmijn Bastings, Ben Poole, Rianne van den Berg, and Tim Salimans. Autoregressive diffusion models. *arXiv preprint arXiv:2110.02037*, 2021a.
- Emiel Hoogetboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in Neural Information Processing Systems*, 34:12454–12465, 2021b.
- Minghui Hu, Yujie Wang, Tat-Jen Cham, Jianfei Yang, and Ponnuthurai N Suganthan. Global context with discrete diffusion in vector quantised modelling for image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11502–11511, 2022.

- Daniel Zhengyu Huang, Jiaoyang Huang, and Zhengjiang Lin. Convergence analysis of probability flow ode for score-based generative models. *arXiv preprint arXiv:2404.09730*, 2024.
- Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- Ilia Igashov, Arne Schneuing, Marwin Segler, Michael Bronstein, and Bruno Correia. Retrobridge: Modeling retrosynthesis with markov bridges. *arXiv preprint arXiv:2308.16212*, 2023.
- Naoto Inoue, Kotaro Kikuchi, Edgar Simo-Serra, Mayu Otani, and Kota Yamaguchi. Layoutdm: Discrete diffusion model for controllable layout generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10167–10176, 2023.
- Jean Jacod and Albert Shiryaev. *Limit theorems for stochastic processes*, volume 288. Springer Science & Business Media, 2013.
- Frank P Kelly. *Reversibility and stochastic networks*. Cambridge University Press, 2011.
- Thomas J Kerby and Kevin R Moon. Training-free guidance for discrete diffusion models for molecular generation. *arXiv preprint arXiv:2409.07359*, 2024.
- Steven P Lalley. Continuous time markov chains. *Lecture Notes, University of Chicago*, pp. 34, 2012.
- Günter Last and Mathew Penrose. *Lectures on the Poisson process*, volume 7. Cambridge University Press, 2017.
- Michel Ledoux. Concentration of measure and logarithmic sobolev inequalities. In *Seminaire de probabilites XXXIII*, pp. 120–216. Springer, 2006.
- Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence for score-based generative modeling with polynomial complexity. *Advances in Neural Information Processing Systems*, 35:22870–22882, 2022.
- Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence of score-based generative modeling for general data distributions. In *International Conference on Algorithmic Learning Theory*, pp. 946–985. PMLR, 2023.
- Tzong-Yow Lee and Horng-Tzer Yau. Logarithmic sobolev inequality for some models of random walks. *The Annals of Probability*, 26(4):1855–1873, 1998.
- Jose Lezama, Tim Salimans, Lu Jiang, Huiwen Chang, Jonathan Ho, and Irfan Essa. Discrete predictor-corrector diffusion models for image synthesis. In *The Eleventh International Conference on Learning Representations*, 2022.
- Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. Towards faster non-asymptotic convergence for diffusion-based generative models. *arXiv preprint arXiv:2306.09251*, 2023.
- Gen Li, Yu Huang, Timofey Efimov, Yuting Wei, Yuejie Chi, and Yuxin Chen. Accelerating convergence of score-based diffusion models, provably. *arXiv preprint arXiv:2403.03852*, 2024a.
- Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. Towards non-asymptotic convergence for diffusion-based generative models. In *The Twelfth International Conference on Learning Representations*, 2024b.
- Gen Li, Yuting Wei, Yuejie Chi, and Yuxin Chen. A sharp convergence theory for the probability flow odes of diffusion models. *arXiv preprint arXiv:2408.02320*, 2024c.
- Tiejun Li. Analysis of explicit tau-leaping schemes for simulating chemically reacting systems. *Multiscale Modeling & Simulation*, 6(2):417–436, 2007.
- Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Biancalani, Aviv Regev, Sergey Levine, and Masatoshi Uehara. Derivative-free guidance in continuous and discrete diffusion models with soft value-based decoding. *arXiv preprint arXiv:2408.08252*, 2024d.

- Yang Li, Jinpei Guo, Runzhong Wang, and Junchi Yan. From distribution learning in training to gradient search in testing for combinatorial optimization. *Advances in Neural Information Processing Systems*, 36, 2024e.
- Yuchen Liang, Peizhong Ju, Yingbin Liang, and Ness Shroff. Non-asymptotic convergence of discrete-time diffusion models: New approach and improved rate. *arXiv preprint arXiv:2402.13901*, 2024.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.
- Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Forty-first International Conference on Machine Learning*, 2024.
- Yunhong Lou, Linchao Zhu, Yaxiong Wang, Xiaohan Wang, and Yi Yang. Diversemotion: Towards diverse human motion generation via discrete diffusion. *arXiv preprint arXiv:2309.01372*, 2023.
- Sokhna Diarra Mbacke and Omar Rivasplata. A note on the convergence of denoising diffusion probabilistic models. *arXiv preprint arXiv:2312.05989*, 2023.
- Chenlin Meng, Kristy Choi, Jiaming Song, and Stefano Ermon. Concrete score matching: Generalized score matching for discrete data. *Advances in Neural Information Processing Systems*, 35: 34532–34545, 2022.
- Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking guidance for discrete state-space diffusion and flow models. *arXiv preprint arXiv:2406.01572*, 2024.
- Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon. Permutation invariant graph generation via score-based generative modeling. In *International Conference on Artificial Intelligence and Statistics*, pp. 4474–4484. PMLR, 2020.
- Bernt Øksendal. *Stochastic differential equations*. Springer, 2003.
- Francesco Pedrotti, Jan Maas, and Marco Mondelli. Improved convergence of score-based diffusion models via prediction-correction. *arXiv preprint arXiv:2305.14164*, 2023.
- Philip Protter. Point process differentials with evolving intensities. In *Nonlinear stochastic problems*, pp. 467–472. Springer, 1983.
- Yiming Qin, Clement Vignac, and Pascal Frossard. Sparse training of discrete diffusion models for graph generation. *arXiv preprint arXiv:2311.02142*, 2023.
- Yinuo Ren, Haoxuan Chen, Yuchen Zhu, Wei Guo, Yongxin Chen, Grant M Rotskoff, Molei Tao, and Lexing Ying. Fast solvers for discrete diffusion models: Theory and applications of high-order algorithms. *arXiv preprint arXiv:2502.00234*, 2025a.
- Yinuo Ren, Grant M Rotskoff, and Lexing Ying. A unified approach to analysis and design of denoising markov models. *arXiv preprint arXiv:2504.01938*, 2025b.
- Pierre H Richemond, Sander Dieleman, and Arnaud Doucet. Categorical sdes with simplex diffusion. *arXiv preprint arXiv:2210.14784*, 2022.
- Eric A Riesel, Tsach Mackey, Hamed Nilforoshan, Minkai Xu, Catherine K Badding, Alison B Altman, Jure Leskovec, and Danna E Freedman. Crystal structure determination from powder diffraction patterns with generative machine learning. *Journal of the American Chemical Society*, 2024.
- Severi Rissanen, Markus Heinonen, and Arno Solin. Improving discrete diffusion models via structured preferential generation. *arXiv preprint arXiv:2405.17889*, 2024.

- Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language models. *arXiv preprint arXiv:2406.07524*, 2024.
- Laurent Saloff-Coste. Lectures on finite markov chains. lectures on probability theory and statistics (saint-flour, 1996), 301–413. *Lecture Notes in Math*, 1665, 1997.
- Javier E Santos, Zachary R Fox, Nicholas Lubbers, and Yen Ting Lin. Blackout diffusion: generative diffusion models in discrete-state spaces. In *International Conference on Machine Learning*, pp. 9034–9059. PMLR, 2023.
- Ari Seff, Wenda Zhou, Farhan Damani, Abigail Doyle, and Ryan P Adams. Discrete object generation with reversible inductive construction. *Advances in neural information processing systems*, 32, 2019.
- Nikil Roashan Selvam, Amil Merchant, and Stefano Ermon. Self-refining diffusion samplers: Enabling parallelization via parareal iterations. *arXiv preprint arXiv:2412.08292*, 2024.
- Chence Shi, Minkai Xu, Zhaocheng Zhu, Weinan Zhang, Ming Zhang, and Jian Tang. Graphaf: a flow-based autoregressive model for molecular graph generation. *arXiv preprint arXiv:2001.09382*, 2020.
- Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis K Titsias. Simplified and generalized masked diffusion for discrete data. *arXiv preprint arXiv:2406.04329*, 2024.
- Andy Shih, Suneel Belkhale, Stefano Ermon, Dorsa Sadigh, and Nima Anari. Parallel sampling of diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pp. 2256–2265. PMLR, 2015.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in neural information processing systems*, 34:1415–1428, 2021.
- Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and Tommi Jaakkola. Dirichlet flow matching with applications to dna sequence design. *arXiv preprint arXiv:2402.05841*, 2024.
- Daniel W Stroock. Logarithmic sobolev inequalities for gibbs states. In *Dirichlet forms*, 1993.
- Haoran Sun, Lijun Yu, Bo Dai, Dale Schuurmans, and Hanjun Dai. Score-based continuous-time discrete diffusion models. *arXiv preprint arXiv:2211.16750*, 2022.
- Zhiwei Tang, Jiasheng Tang, Hao Luo, Fan Wang, and Tsung-Hui Chang. Accelerating parallel sampling of diffusion models. In *Forty-first International Conference on Machine Learning*, 2024.
- Belinda Tzen and Maxim Raginsky. Theoretical guarantees for sampling and inference in generative models with latent diffusions. In *Conference on Learning Theory*, pp. 3084–3114. PMLR, 2019.
- Nico M Van Dijk. Uniformization for nonhomogeneous markov chains. *Operations research letters*, 12(5):283–291, 1992.
- Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. *arXiv preprint arXiv:2209.14734*, 2022.

- Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- Lingxiao Wang, Gert Aarts, and Kai Zhou. Generative diffusion models for lattice field theory. *arXiv preprint arXiv:2311.03578*, 2023.
- Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eisenach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.
- Ludwig Winkler, Lorenz Richter, and Manfred Opperr. Bridging discrete and continuous state spaces: Exploring the ehrenfest process in time-continuous diffusion models. *arXiv preprint arXiv:2405.03549*, 2024.
- Tong Wu, Zhihao Fan, Xiao Liu, Hai-Tao Zheng, Yeyun Gong, Jian Jiao, Juntao Li, Jian Guo, Nan Duan, Weizhu Chen, et al. Ar-diffusion: Auto-regressive diffusion model for text generation. *Advances in Neural Information Processing Systems*, 36:39957–39974, 2023.
- Zhichao Wu, Qiulin Li, Sixing Liu, and Qun Yang. Dctts: Discrete diffusion model with contrastive learning for text-to-speech generation. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 11336–11340. IEEE, 2024.
- Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*, 2022.
- Dongchao Yang, Jianwei Yu, Helin Wang, Wen Wang, Chao Weng, Yuexian Zou, and Dong Yu. Diffsound: Discrete diffusion model for text-to-sound generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:1720–1733, 2023a.
- John J Yang, Jason Yim, Regina Barzilay, and Tommi Jaakkola. Fast non-autoregressive inverse folding with discrete diffusion. *arXiv preprint arXiv:2312.02447*, 2023b.
- Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023c.
- Kai Yi, Bingxin Zhou, Yiqing Shen, Pietro Liò, and Yuguang Wang. Graph denoising diffusion for inverse protein folding. *Advances in Neural Information Processing Systems*, 36, 2024.
- Junyi Zhang, Jiaqi Guo, Shizhao Sun, Jian-Guang Lou, and Dongmei Zhang. Layoutdiffusion: Improving graphic layout generation by discrete diffusion probabilistic models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7226–7236, 2023.
- Linfeng Zhang, Weinan E, and Lei Wang. Monge-ampère flow for generative modeling. *arXiv preprint arXiv:1809.10188*, 2018.
- Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. *arXiv preprint arXiv:2204.13902*, 2022.
- Zikun Zhang, Zixiang Chen, and Quanquan Gu. Convergence of score-based discrete diffusion models: A discrete-time analysis. *arXiv preprint arXiv:2410.02321*, 2024.
- Lingxiao Zhao, Xueying Ding, Lijun Yu, and Leman Akoglu. Improving and unifying discrete&continuous-time discrete denoising diffusion. *arXiv preprint arXiv:2402.03701*, 2024.
- Lin Zheng, Jianbo Yuan, Lei Yu, and Lingpeng Kong. A reparameterized discrete diffusion model for text generation. *arXiv preprint arXiv:2302.05737*, 2023.
- Kun Zhou, Yifan Li, Wayne Xin Zhao, and Ji-Rong Wen. Diffusion-nat: Self-prompting discrete diffusion for non-autoregressive text generation. *arXiv preprint arXiv:2305.04044*, 2023.
- Ye Zhu, Yu Wu, Kyle Olszewski, Jian Ren, Sergey Tulyakov, and Yan Yan. Discrete contrastive diffusion for cross-modal music and image generation. *arXiv preprint arXiv:2206.07771*, 2022.
- Yuchen Zhu, Tianrong Chen, Evangelos A Theodorou, Xie Chen, and Molei Tao. Quantum state generation with structure-preserving diffusion model. *arXiv preprint arXiv:2404.06336*, 2024.

A MATHEMATICAL FRAMEWORK OF POISSON RANDOM MEASURE

In this section, we provide a mathematical framework for Poisson random measure with evolving intensity, which is crucial for the error analysis of discrete diffusion models in the main text.

A.1 PRELIMINARIES

We first provide the definition of the ordinary Poisson random measure.

Definition A.1 (Poisson Random Measure). *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and $(\mathbb{X}, \mathcal{B}, \nu)$ be a measure space satisfying that*

$$\int_{\mathbb{X}} 1 \vee |y| \vee |y|^2 \nu(dy) < \infty,$$

The random measure $N(dt, dy)$ on $\mathbb{R}^+ \times \mathbb{X}$ is called a Poisson random measure w.r.t. measure ν if it is a random counting measure satisfying the following properties:

- (i) *For any $B \in \mathcal{B}$ and $0 \leq s < t$, $N((s, t] \times B) \sim \mathcal{P}(\nu(B)(t - s))$;*
- (ii) *For any $t \geq 0$ and pairwise disjoint sets $\{B_i\}_{i \in [n]} \subset \mathcal{B}$, $\{N_t(B_i) := N((0, t] \times B_i)\}_{i \in [n]}$ are independent stochastic processes.*

The following definition of *predictability* will be frequently used for the well-definedness of stochastic integrals w.r.t. Poisson random measure, and thus the extension from ordinary Poisson random measure to Poisson random measure with evolving intensity.

Definition A.2 (Predictability). *The predictable σ -algebra on $\mathbb{R}^+ \times \mathbb{X}$ is defined as the σ -algebra generated by all sets of the form $(s, t] \times B$ for $0 \leq s < t$ and $B \in \mathcal{B}$. A process X_t is called predictable if and only if X_t is predictable w.r.t. the predictable σ -algebra above.*

In the following, we will define the Poisson random measure with evolving intensity, which is a special case of random measures (Jacod & Shiryaev, 2013, Definition 1.3).

Definition A.3 (Poisson Random Measure with Evolving Intensity). *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and $(\mathbb{X}, \mathcal{B}, \nu)$ be a measure space. Suppose $\lambda_t(y)$ is a non-negative predictable process on $\mathbb{R}^+ \times \mathbb{X} \times \Omega$ satisfying that for any $0 \leq T < \bar{T}$,*

$$\int_0^T \int_{\mathbb{X}} 1 \vee |y| \vee |y|^2 \lambda_t(y) \nu(dy) dt < \infty, \text{ a.s..}$$

The random measure $N[\lambda](dt, dy)$ on $\mathbb{R}^+ \times \mathbb{X}$ is called a Poisson random measure with evolving intensity $\lambda_t(y)$ w.r.t. measure ν if it is a random counting measure satisfying the following properties:

- (i) *For any $B \in \mathcal{B}$ and $0 \leq s < t$, $N[\lambda]((s, t] \times B) \sim \mathcal{P}\left(\int_s^t \int_B \lambda_\tau(y) \nu(dy) d\tau\right)$;*
- (ii) *For any $t \geq 0$ and pairwise disjoint sets $\{B_i\}_{i \in [n]} \subset \mathcal{B}$,*

$$\{N_t[\lambda](B_i) := N[\lambda]((0, t] \times B_i)\}_{i \in [n]}$$

are independent stochastic processes.

Theorem A.4 (Well-definedness of Poisson Random Measure with Evolving Intensity). *The Poisson random measure $N[\lambda](dt, dy)$ with evolving intensity $\lambda_t(y)$ is well-defined under the conditions in the definition above.*

Proof. We first augment the $(\mathbb{X}, \mathcal{B}, \nu)$ measure space to a product space $(\mathbb{X} \times \mathbb{R}, \mathcal{B} \times \mathcal{B}(\mathbb{R}), \nu \times m)$, where m is the Lebesgue measure on \mathbb{R} , and $\mathcal{B}(\mathbb{R})$ is the Borel σ -algebra on \mathbb{R} . The Poisson random measure with evolving intensity $\lambda_t(y)$ can be defined in the augmented measure space as

$$N[\lambda]((s, t] \times B) := \int_s^t \int_B \int_{\mathbb{R}} \mathbf{1}_{0 \leq \xi \leq \lambda_\tau(y)} N(d\tau, dy, d\xi), \quad (\text{A.1})$$

where $N(d\tau, dy, d\xi)$ is the Poisson random measure on $\mathbb{R}^+ \times \mathbb{X} \times \mathbb{R}$ w.r.t. measure $\nu(dy)d\xi$.

Then it is straightforward to verify the two conditions in the definition of Poisson random measure with evolving intensity by noticing that for pairwise disjoint sets $\{B_i\}_{i \in [n]} \subset \mathcal{B}$, $\{B_i \times \mathbb{R}\}_{i \in [n]} \subset \mathcal{B} \times \mathcal{B}(\mathbb{R})$ are also pairwise disjoint.

The Poisson random process $N[\lambda](dt, dy)$ with evolving intensity $\lambda_t(y)$ is well-defined up to an eventual explosion time

$$\bar{T} = \inf_T \left\{ \int_0^T \int_{\mathbb{X}} \lambda_t(y) \nu(dy) dt = \infty, \text{ a.s.} \right\}.$$

We refer the readers to (Protter, 1983) for a more rigorous detailed version of the proof. \square

Remark A.5 (Relations to the classical Poisson random measure and that with state-dependent density). *The classical Poisson random measure is well-studied by the theory of Lévy processes (Applebaum, 2009), and the extension to the state-dependent intensity is proposed and analyzed in (Glasserman & Merener, 2004). Notably, Li (2007) establishes the stochastic integral formulation for the chemical master equation with the Poisson random measure with state-dependent intensity, which is a special case of the evolving intensity, and subsequently shows the weak and strong convergence of the τ -leaping scheme.*

Remark A.6 (Relation to the Cox process). *The Poisson random measure with evolving intensity shares multiple similarities with the Cox process (Cox, 1955; Last & Penrose, 2017), including being a point process and with the intensity being a random measure. The main difference is that the Cox process is defined on a general measure space, while the Poisson random measure with evolving intensity is defined on the product space $(\mathbb{X} \times \mathbb{R}, \mathcal{B} \times \mathcal{B}(\mathbb{R}), \nu \times m)$ and the intensity function is required to be predictable to ensure the well-definedness of its stochastic integral.*

A.2 STOCHASTIC INTEGRAL W.R.T. POISSON RANDOM MEASURE

The following theorems provide the properties of stochastic integrals w.r.t. Poisson random measure with evolving intensity. The proofs are based on the observation that with the augmentation of the measure space argument (A.1), the stochastic integral w.r.t. Poisson random measure with evolving intensity in $(\mathbb{X}, \mathcal{B}, \nu)$ can be reduced to the stochastic integral w.r.t. homogeneous Poisson random measure in $(\mathbb{X} \times \mathbb{R}, \mathcal{B} \times \mathcal{B}(\mathbb{R}), \nu \times m)$, and under certain conditions on the measure space $(\mathbb{X}, \mathcal{B}, \nu)$, to the well-known Lévy-type stochastic integral (Applebaum, 2009). For simplicity, we will work on the interval $t \in [0, T]$ with $T < \bar{T}$ and the following regularity conditions of the Poisson random measure:

$$0 < \operatorname{ess\,inf}_{\tau \in [0, T], y \in \mathbb{X}} \lambda_\tau(y) \leq \operatorname{ess\,sup}_{\tau \in [0, T], y \in \mathbb{X}} \lambda_\tau(y) < +\infty.$$

One can easily generalize the following results to their local versions on $[0, \bar{T})$ by considering its compact subsets.

Theorem A.7 (Stochastic Integrals w.r.t. Poisson Random Measure with Evolving Density). *For any predictable process $K_t(y)$ on $\mathbb{R}^+ \times \mathbb{X} \times \Omega$, the stochastic integral w.r.t. Poisson random measure with evolving intensity $\lambda_t(y)$*

$$x_t = x_0 + \int_0^t \int_{\mathbb{X}} K_\tau(y) N[\lambda](d\tau, dy), \quad (\text{A.2})$$

has a unique solution, for which the following properties hold:

(1) (Expectation) *For any $t \geq 0$, we have*

$$\mathbb{E} \left[\int_0^t \int_{\mathbb{X}} K_\tau(y) N[\lambda](d\tau, dy) \right] = \int_0^t \int_{\mathbb{X}} K_\tau(y) \lambda_\tau(y) \nu(dy) d\tau;$$

(2) (Martingale) *For any $t \geq 0$, we have*

$$\int_0^t \int_{\mathbb{X}} K_\tau(y) \tilde{N}[\lambda](d\tau, dy) := \int_0^t \int_{\mathbb{X}} K_\tau(y) N[\lambda](d\tau, dy) - \int_0^t \int_{\mathbb{X}} K_\tau(y) \lambda_\tau(y) \nu(dy) d\tau$$

is a local \mathcal{F}_t -martingale;

(3) (*Itô Isometry*) For any $t \geq 0$, we have

$$\mathbb{E} \left[\left(\int_0^t \int_{\mathbb{X}} K_{\tau}(y) N[\lambda](d\tau, dy) \right)^2 \right] = \int_0^t \int_{\mathbb{X}} K_{\tau}(y)^2 \lambda_{\tau}(y) \nu(dy) d\tau.$$

Proof. We first write the integral (A.2) in the augmented measure space $(\mathbb{X} \times \mathbb{R}, \mathcal{B} \times \mathcal{B}(\mathbb{R}), \nu \times m)$ as

$$x_t = x_0 + \int_0^t \int_{\mathbb{X}} \int_{\mathbb{R}} K_{\tau}(y) \mathbf{1}_{0 \leq \xi \leq \lambda_{\tau}(y)} N(d\tau, dy, d\xi), \quad (\text{A.3})$$

and since $K_t(y) \mathbf{1}_{0 \leq \xi \leq \lambda_t(y)}$ is a predictable process, the desired properties can be derived from the corresponding properties of the stochastic integral w.r.t. Poisson random measure in the augmented measure space.

The subsequent proof will follow a similar argument as the proof of the stochastic integral w.r.t. Brownian motion (e.g. in (Øksendal, 2003)) by starting from proving the properties for elementary processes, which in our case refer to working with the *elementary predictable processes* of the following form:

$$Z_t(y, \xi)(\omega) = \sum_{i=0}^{n-1} \sum_{j=1}^m \sum_{k=1}^l Z_{i,j,k}(\omega) \mathbf{1}_{t \in (t_i, t_{i+1}]} \mathbf{1}_{y \in B_j} \mathbf{1}_{\xi \in C_k},$$

where $0 = t_0 < \dots < t_n = T$ is a partition of $[0, T]$, $B_j \in \mathcal{B}$ for $j \in [m]$ are a partition of \mathbb{X} with $\nu(B_k) < \infty$, and $C_k \in \mathcal{B}(\mathbb{R})$ for $k \in [l]$ are a partition of the time interval $[0, \text{ess sup}_{\tau \in [0, T], y \in \mathbb{X}} \lambda_{\tau}(y)]$ with $m(C_k) < \infty$, and $Z_{i,j,k}$ is bounded and \mathcal{F}_{t_i} -measurable, on which the stochastic integral is defined as

$$\int_0^t \int_{\mathbb{X}} Z_{\tau}(y, \xi) N(d\tau, dy, d\xi) = \sum_{i=0}^{n-1} \sum_{j=1}^m \sum_{k=1}^l Z_{i,j,k} N_{t_i}((t_i, t_{i+1}] \times B_j \times C_k).$$

Then, it is straightforward to verify the properties of the stochastic integral for the elementary predictable process $Z_t^+(y, \xi)$, using the definition of Poisson random measure (Definition A.1). For general predictable processes $Z_t(y, \xi)$, we write $Z_t(y, \xi) = Z_t^+(y, \xi) - Z_t^-(y, \xi)$, where $Z_t^+(y, \xi)$ and $Z_t^-(y, \xi)$ are positive and negative parts of $Z_t(y, \xi)$, and apply the results to $Z_t^+(y, \xi)$ and $Z_t^-(y, \xi)$ separately.

Finally, we take $Z_t(y, \xi) = K_t(y) \mathbf{1}_{0 \leq \xi \leq \lambda_t(y)}$ to derive the properties of the stochastic integral w.r.t. Poisson random measure with evolving intensity.

We refer readers to (Eberle, 2015, Section 2.2) for detailed arguments. For the uniqueness of the solution to the stochastic integral, we also refer to (Protter, 1983, Theorem 3.1). \square

Proposition A.8. Define the list of jump times $(t_n)_{n \in \mathbb{N}}$ recursively as

$$t_0 = 0, \quad t_{n+1} = \inf \{t > t_n | \Delta x_t \neq 0\}, \quad n \geq 0,$$

the Poisson random measure $N[\lambda](dt, dy)$ with evolving intensity $\lambda_t(y)$ can be written as

$$N[\lambda](dt, dy) = \sum_{n=1}^{\infty} \delta_{t_n}(dt) \delta_{Y_n}(dy), \quad (\text{A.4})$$

and the stochastic integral (A.2) is càdlàg and can be rewritten as a sum of jumps:

$$x_t = x_0 + \sum_{n=1}^N \Delta x_{t_n} = x_0 + \sum_{n=1}^N K_{t_n}(Y_n), \quad (\text{A.5})$$

where N is a random variable satisfying $t_N \leq t < t_{N+1}$, and Δx_{t_n} are the jumps $\Delta x_{t_n} = x_{t_n} - x_{t_n}^-$ with $x_{t_n}^- := \lim_{s \rightarrow t_n^-} x_s$.

Proof. To see the solution is càdlàg, we notice the following right limit at time t :

$$\lim_{\epsilon \rightarrow 0} (x_{t+\epsilon} - x_t) = \int_{(t, t+\epsilon] \times \mathbb{X}} K_t(y) N[\lambda](dt, dy) \rightarrow 0,$$

and the left limit at time t :

$$\Delta x_t = \lim_{\epsilon \rightarrow 0} (x_t - x_{t-\epsilon}) = \int_{(t-\epsilon, t] \times \mathbb{X}} K_t(y) N[\lambda](dt, dy) \rightarrow \int_{\mathbb{X}} K_t(y) N[\lambda](\{t\} \times dy), \quad (\text{A.6})$$

where the notation $N[\lambda](\{t\} \times dy)$ should be understood as $N[\lambda](\{t\} \times dy) = 0$ if $t \notin \{t_n\}_{n \in \mathbb{N}}$, or otherwise $Y_n = N[\lambda](\{t_n\} \times dy)$ is a random variable on \mathbb{X} .

Since the Poisson random measure $N[\lambda](dt, dy)$ with evolving intensity $\lambda_t(y)$ is a random counting measure, it can be represented as a countable sum of Dirac measures as in (A.4), and thus we have

$$\begin{aligned} x_t &= x_0 + \int_0^t \int_{\mathbb{X}} K_t(y) N[\lambda](dt, dy) \\ &= x_0 + \int_0^t \int_{\mathbb{X}} K_t(y) \sum_{n=1}^N \delta_{t_n}(dt) \delta_{Y_n}(dy) = x_0 + \sum_{n=1}^N K_{t_n}(Y_n). \end{aligned}$$

By the definition of Poisson random measure, $(t_n)_{n \in [N]}$ are also the jump times of the homogeneous Poisson random measure $N(dt, dy, d\xi)$ in the augmented measure space $(\mathbb{X} \times \mathbb{R}, \mathcal{B} \times \mathcal{B}(\mathbb{R}), \nu \times m)$ w.r.t. measure $\nu(dy)d\xi$. Therefore, with a slight abuse of notations, we will assume $(t_n, Y_n, \Xi_n)_{n \in [N]}$ are i.i.d. random variables with probability measure proportional to $dt\nu(dy)d\xi$, for each of which $\Xi_n \leq \lambda_{t_n}(Y_n)$ holds because otherwise the jump would not occur.

Then the distribution of Y_n can be derived as a conditional probability of the jump location Y_n given the jump time t_n and $\Xi_n \leq \lambda_{t_n}(Y_n)$:

$$\mathbb{P}(Y_n = y) \nu(dy) = \frac{\int_{\mathbb{R}} \nu(dy) \mathbf{1}_{\Xi_n \leq \lambda_{t_n}(y)} d\xi}{\int_{\mathbb{R}} \int_{\mathbb{X}} \nu(dy) \mathbf{1}_{\Xi_n \leq \lambda_{t_n}(y)} d\xi} = \frac{\lambda_{t_n}(y) \nu(dy)}{\int_{\mathbb{X}} \lambda_{t_n}(y) \nu(dy)}, \quad (\text{A.7})$$

and the proof is complete. \square

The following theorem gives the martingale characterization of Poisson random measure with evolving intensity, which will be crucial for the proof of the change of measure arguments:

Theorem A.9 (Martingale Characterization of Poisson Random Measure with Evolving Density). *Let $N[\lambda](dt, dy)$ be a \mathcal{F}_t -adapted process in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Then $N[\lambda](dt, dy)$ is a Poisson random measure with evolving intensity $\lambda_t(y)$ if and only if the complex-valued process*

$$M_t[f] = \exp \left(i \int_0^t \int_{\mathbb{X}} f_\tau(y) N[\lambda](d\tau, dy) + \int_0^t \int_{\mathbb{X}} \left(1 - e^{if_\tau(y)} \right) \lambda_\tau(y) \nu(dy) d\tau \right) \quad (\text{A.8})$$

is a local martingale for any predictable process $f_\tau(y)$ satisfying that $f_\tau(y) \in L^1(\mathbb{X}, \nu)$, a.s..

Proof. By Proposition A.8, we rewrite the stochastic integral as a sum of jumps:

$$\int_0^t \int_{\mathbb{X}} f_t(y) N[\lambda](d\tau, dy) = \sum_{n=1}^N f_{t_n}(Y_n),$$

where $(t_n, Y_n, \Xi_n)_{n \in [N]}$ are i.i.d. random variables with probability measure proportional to $dt\nu(dy)d\xi$, for each of which $\Xi_n \leq \lambda_{t_n}(Y_n)$ holds, following a similar argument as in the proof of Proposition A.8.

Then, it is straightforward to derive the following probability of the jump time $t_n = \tau$:

$$\mathbb{P}(t_n = \tau) d\tau = \frac{\int_{\mathbb{X}} \mathbb{P}(Y_n = y, t_n = \tau) \nu(dy) d\tau}{\int_0^t \int_{\mathbb{X}} \mathbb{P}(Y_n = y, t_n = \tau) \nu(dy) d\tau} = \frac{\int_{\mathbb{X}} \lambda_\tau(y) \nu(dy) d\tau}{\int_0^t \int_{\mathbb{X}} \lambda_\tau(y) \nu(dy) d\tau};$$

and by the definition of the Poisson random measure, we have the following probability of the total number of jumps $N = n$:

$$\mathbb{P}(N = n) = \frac{1}{n!} \exp \left(- \int_0^t \int_{\mathbb{X}} \lambda_{\tau}(y) \nu(dy) d\tau \right) \left(\int_0^t \int_{\mathbb{X}} \lambda_{\tau}(y) \nu(dy) d\tau \right)^n.$$

Without loss of generality, we only verify $\mathbb{E}[M_t[f]] = 1$ as follows, and general cases are similar by Markov property:

$$\begin{aligned} & \mathbb{E} \left[\exp \left(i \int_0^t \int_{\mathbb{X}} f_{t_n}(y) N[\lambda](d\tau, dy) \right) \right] \\ &= \mathbb{E} \left[\exp \left(i \sum_{n=1}^N f_{t_n}(Y_n) \right) \middle| \Xi_n \leq \lambda_{t_n}(Y_n), \forall n \in [N] \right] \\ &= \mathbb{E} \left[\prod_{n=1}^N \mathbb{E} \left[e^{i f_{t_n}(Y_n)} \middle| \Xi_n \leq \lambda_{t_n}(Y_n) \right] \right] \\ &= \mathbb{E} \left[\prod_{n=1}^N \mathbb{E} \left[\int_{\mathbb{X}} \frac{e^{i f_{\tau}(y)} \lambda_{\tau}(y) \nu(dy)}{\int_{\mathbb{X}} \lambda_{\tau}(y) \nu(dy)} \middle| t_n = \tau \right] \right] \\ &= \sum_{n=1}^{\infty} \frac{1}{n!} \left(\int_0^t \int_{\mathbb{X}} e^{i f_{\tau}(y)} \lambda_{\tau}(y) \nu(dy) d\tau \right)^n \exp \left(- \int_0^t \int_{\mathbb{X}} \lambda_{\tau}(y) \nu(dy) d\tau \right) \\ &= \exp \left(\int_0^t \int_{\mathbb{X}} (e^{i f_{\tau}(y)} - 1) \lambda_{\tau}(y) \nu(dy) d\tau \right), \end{aligned}$$

which immediately yields the desired result $\mathbb{E}[M_t[f]] = 1$.

On the other hand, for any $0 \leq s < t$ and $B \in \mathcal{B}$, we set

$$Z_t(y) = u \mathbf{1}_{t \in (s, t]} \mathbf{1}_{y \in B},$$

where $u \in \mathbb{R}$, and by assumption, we have

$$\begin{aligned} & \mathbb{E}[M_t[Z]] = \mathbb{E} \left[\exp \left(i \int_0^t \int_{\mathbb{X}} Z_{\tau}(y) N[\lambda](d\tau, dy) + \int_0^t \int_{\mathbb{X}} (1 - e^{i Z_{\tau}(y)}) \lambda_{\tau}(y) \nu(dy) d\tau \right) \right] \\ &= \mathbb{E} \left[\exp \left(i u \int_s^t \int_B N[\lambda](d\tau, dy) + \int_s^t \int_{\mathbb{X}} (1 - e^{iu}) \lambda_{\tau}(y) \nu(dy) d\tau \right) \right] \\ &= \mathbb{E} \left[\exp \left(i u N[\lambda]((s, t] \times B) + (1 - e^{iu}) (t - s) \int_{\mathbb{X}} \lambda_{\tau}(y) \nu(dy) \right) \right] = 1, \end{aligned}$$

i.e. the following holds for any $u \in \mathbb{R}$:

$$\mathbb{E}[\exp(iu N[\lambda]((s, t] \times B))] = (e^{iu} - 1) (t - s) \int_{\mathbb{X}} \lambda_{\tau}(y) \nu(dy),$$

which by Lévy's continuity theorem implies that

$$N[\lambda]((s, t] \times B) \sim \mathcal{P} \left((t - s) \int_{\mathbb{X}} \lambda_{\tau}(y) \nu(dy) \right),$$

and thus $N[\lambda](dt, dy)$ is a Poisson random measure with evolving intensity $\lambda_t(y)$ by Definition A.1. \square

Theorem A.10 (Itô's Formula for Poisson Random Measure with Evolving Density). *Let $N[\lambda](dt, dy)$ be a Poisson random measure with evolving intensity $\lambda_t(y)$ in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and $K_t(y)$ be a predictable process on $\mathbb{R}^+ \times \mathbb{X} \times \Omega$. Suppose a process x_t satisfies the stochastic integral*

$$x_t = x_0 + \int_0^t \int_{\mathbb{X}} K_{\tau}(y) N[\lambda](d\tau, dy), \quad (\text{A.9})$$

then for any $f_t(y) \in C(\mathbb{R}^+ \times \mathbb{X})$ with probability 1, we have

$$f_t(x_t) = f_0(x_0) + \int_0^t \partial_\tau f_\tau(x_\tau) d\tau + \int_0^t \int_{\mathbb{X}} (f_\tau(x_{\tau-} + K_\tau(y)) - f_\tau(x_{\tau-})) N[\lambda](d\tau, dy).$$

Proof. By Proposition A.8, we again rewrite the stochastic integral as a sum of jumps:

$$x_t = x_0 + \sum_{n=1}^N K_{t_n}(Y_n),$$

where $(t_n)_{n \in [N]}$ are the jump times with $t_0 = 0$ and $t_N \leq t < t_{N+1}$, and $(Y_n)_{n \in [N]}$ are the jump locations. Consequently, it is easy to see that $x_{t_n^-} = x_t = x_{t_{n-1}}$ for $t \in (t_{n-1}, t_n]$ and $n \in [N]$.

Then we have the following decomposition:

$$\begin{aligned} & f_t(x_t) - f_0(x_0) \\ &= f_t(x_t) - f_{t_N}(x_{t_N}) + \sum_{n=1}^N \left(f_{t_n}(x_{t_n}) - f_{t_n}(x_{t_n^-}) + f_{t_n}(x_{t_n^-}) - f_{t_{n-1}}(x_{t_{n-1}}) \right) \\ &= \int_{t_N}^t \partial_\tau f_\tau(x_{t_N}) d\tau + \sum_{n=1}^N \int_{t_{n-1}}^{t_n} \partial_\tau f_\tau(x_{t_{n-1}}) d\tau + \sum_{n=1}^N (f_{t_n}(x_{t_{n-1}} + K_{t_n}(Y_n)) - f_{t_n}(x_{t_{n-1}})), \end{aligned}$$

and for the last term in the above equation, we have

$$\begin{aligned} & \sum_{n=1}^N (f_{t_n}(x_{t_{n-1}} + K_{t_n}(Y_n)) - f_{t_n}(x_{t_{n-1}})) = \sum_{n=1}^N (f_{t_n}(x_{t_n^-} + K_{t_n}(Y_n)) - f_{t_n}(x_{t_n^-})) \\ &= \int_0^t \int_{\mathbb{X}} (f_\tau(x_{\tau-} + K_\tau(y)) - f_\tau(x_{\tau-})) N[\lambda](d\tau, dy). \end{aligned}$$

Combining the above results, we have the desired result. \square

Lemma A.11. Denote the trajectory obtained by simulating the master equation (2.7) of the forward process of the discrete diffusion model as x_t , then the time interval $\Delta t_n = t_{n+1} - t_n$ is distributed according to the following distribution:

$$\mathbb{P}(\Delta t_n > \tau) = \exp \left(\int_0^\tau Q_{\tau'}(x_{t_n}, x_{t_n}) d\tau' \right), \quad (\text{A.10})$$

and the jump location $x_{t_{n+1}}$ is distributed according to the following distribution:

$$\mathbb{P}(x_{t_{n+1}} = y) = - \frac{Q_{t_{n+1}}(y, x_{t_n})}{Q_{t_{n+1}}(x_{t_n}, x_{t_n})}. \quad (\text{A.11})$$

Proof. The results can be found in (Eberle, 2009, Section 1.2). While a fully rigorous proof can be conducted by discretizing the time-inhomogeneous continuous-time Markov chain into 2^n uniform steps and taking the limit as $n \rightarrow \infty$, following the approach in (Lalley, 2012), we will provide a more intuitive proof here for completeness.

Set $p_{t_n} = e_{x_{t_n}}$, where e_y is the y -th unit vector in $\mathbb{R}^{|\mathbb{X}|}$. then the x_{t_n} -th entry of (2.7) yields

$$\frac{d}{dt} \mathbb{P}(x_t = x_{t_n}) = \frac{d}{dt} p_t(x_{t_n}) = \sum_{y \in \mathbb{X}} Q_t(x_{t_n}, y) p_t(y),$$

which, by the assumed continuity of the rate matrix Q_t , implies

$$\mathbb{P}(\Delta t_n > \tau) = \mathbb{P}(x_{t_n+\tau} = x_{t_n}) = 1 + Q_{t_n}(x_{t_n}, x_{t_n})\tau + o(\tau),$$

and thus

$$\frac{d}{d\tau} \log \mathbb{P}(\Delta t_n > \tau) = \lim_{\tau \rightarrow 0} \frac{\log \mathbb{P}(\Delta t_n > \tau)}{\tau} = \lim_{\tau \rightarrow 0} Q_{t_n}(x_{t_n}, x_{t_n}) + o(1) = Q_{t_n}(x_{t_n}, x_{t_n}),$$

integrating the above equation yields the desired result.

Similarly, by setting $\mathbf{p}_{t_{n+1}-\tau} = \mathbf{e}_{x_{t_n}}$, we have for all $y \in \mathbb{X} \setminus \{x_{t_n}\}$:

$$\begin{aligned} \mathbb{P}(x_{t_{n+1}} = y) &= \mathbf{p}_{t_{n+1}}(y) \\ &= \lim_{\tau \rightarrow 0} \frac{\mathbf{p}_{t_{n+1}-\tau}(y) + Q_{t_{n+1}-\tau}(y, x_{t_n})\tau + o(\tau)}{\sum_{y \in \mathbb{X} \setminus \{x_{t_n}\}} (\mathbf{p}_{t_{n+1}-\tau}(y) + Q_{t_{n+1}-\tau}(y, x_{t_n})\tau + o(\tau))} \\ &= \frac{Q_{t_{n+1}}(y, x_{t_n})}{\sum_{y \in \mathbb{X} \setminus \{x_{t_n}\}} Q_{t_{n+1}}(y, x_{t_n})} = -\frac{Q_{t_{n+1}}(y, x_{t_n})}{Q_{t_{n+1}}(x_{t_n}, x_{t_n})}, \end{aligned}$$

and the result follows. \square

A.3 PROOFS OF CHANGE OF MEASURE RELATED ARGUMENTS

Proof of Theorem 3.3. In the following, we will denote the expectation under the measure \mathbb{P} by $\mathbb{E}_{\mathbb{P}}$ and the expectation under the measure \mathbb{Q} by $\mathbb{E}_{\mathbb{Q}}$.

By Theorem A.9, to verify that the Poisson random measure $N[\lambda](dt, dy)$ with evolving intensity $\lambda_t(y)$ is a Poisson random measure with evolving intensity $\lambda_t(y)h_t(y)$ under the measure \mathbb{Q} , it suffices to show that for any $f \in L^1(\mathbb{X}, \nu)$, the complex-valued process

$$M_t[f] = \exp \left(i \int_0^t \int_{\mathbb{X}} f(y) N[\lambda](d\tau, dy) + \int_0^t \int_{\mathbb{X}} (1 - e^{if(y)}) \lambda_{\tau}(y) h_{\tau}(y) \nu(dy) d\tau \right)$$

is a local martingale under the measure \mathbb{Q} .

To this end, we perform the following calculation:

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}}[M_t[f]] &= \mathbb{E}_{\mathbb{P}}[M_t[f] Z_t[h]] \\ &= \mathbb{E}_{\mathbb{P}} \left[\exp \left(i \int_0^t \int_{\mathbb{X}} f(y) N[\lambda](d\tau, dy) + \int_0^t \int_{\mathbb{X}} (1 - e^{if(y)}) \lambda_{\tau}(y) h_{\tau}(y) \nu(dy) d\tau \right) \right. \\ &\quad \left. \exp \left(\int_0^t \int_{\mathbb{X}} \log h_t(y) N[\lambda](dt \times dy) - \int_0^t \int_{\mathbb{X}} (h_t(y) - 1) \lambda_t(y) \nu(dy) \right) \right] \\ &= \mathbb{E}_{\mathbb{P}} \left[\exp \left(i \int_0^t \int_{\mathbb{X}} (f(y) + \log h_t(y)) N[\lambda](d\tau, dy) + \int_0^t \int_{\mathbb{X}} (1 - e^{if(y)h_t(y)}) \lambda_t(y) \nu(dy) \right) \right] \\ &= \mathbb{E}_{\mathbb{P}}[M_t[f + \log h]], \end{aligned}$$

and by assumption, $f + \log h \in L^1(\mathbb{X}, \nu)$, a.s., which implies that $M_t[f + \log h]$ is a local martingale under the measure \mathbb{P} again by Theorem A.9. Consequently, $M_t[f]$ is a local martingale under the measure \mathbb{Q} , and the result follows. \square

Remark A.12. One may verify that the exponential process Z_t in (3.3) is a local \mathcal{F}_t -martingale by Applebaum (2009, Corollary 5.2.2) under mild assumptions on the function $h_t(y)$ (cf. Novikov’s condition in the Girsanov’s theorem for Itô integrals).

Proof of Corollary 3.4. With a similar argument as in the proof of Proposition 3.2, we have the following stochastic integral representation of the approximate backward process with the learned neural network score function \hat{s}_t^{θ} :

$$y_s = y_0 + \int_0^s \int_{\mathbb{X}} (y - y_{s-}) N[\hat{\mu}^{\theta}](ds, dy), \text{ with } \hat{\mu}_s^{\theta}(y) = \tilde{s}_s^{\theta}(y_{s-}, y) \tilde{Q}_s(y_{s-}, y) = \hat{\mu}_s^{\theta}(y). \quad (\text{A.12})$$

By the data-processing inequality and the chain rule of KL divergence, we have

$$D_{\text{KL}}(\tilde{p}_T \| q_T) \leq D_{\text{KL}}(\tilde{p}_{0:T} \| q_{0:T}) = D_{\text{KL}}(\tilde{p}_0 \| q_0) + \mathbb{E}[D_{\text{KL}}(\tilde{p}_{0:T} \| q_{0:T} | \tilde{x}_0 = y_0 = y)].$$

Then notice that conditioning on the alignment of the initial state $\tilde{x}_0 = y_0 = y$ for any $y \in \mathbb{X}$, the second term in the above equation can be expressed as

$$D_{\text{KL}}(\tilde{p}_{0:T} \| q_{0:T} | \tilde{x}_0 = y_0 = y) = \mathbb{E} \left[\log \frac{d\tilde{p}_{0:T}}{dq_{0:T}} \middle| \tilde{x}_0 = y_0 = y \right] = \mathbb{E} \left[\log Z_T^{-1} \left[\frac{\hat{\mu}^\theta}{\mu} \right] \right],$$

where the last equality is by the change of measure in Theorem 3.3 from the stochastic integral formulation (3.2) of the backward process (2.8) with the true score function \tilde{s} to the stochastic integral formulation (A.12) of the approximate backward process with the learned score function \tilde{s}^θ .

Plug in the expression of Z_T in (3.3) and notice that

$$\frac{\hat{\mu}_s^\theta}{\mu_s} = \frac{\tilde{s}_s^\theta(y_{s-}, y) \tilde{Q}_s(y_{s-}, y)}{\tilde{s}_s(y_{s-}, y) \tilde{Q}_s(y_{s-}, y)} = \frac{\tilde{s}_s^\theta(y_{s-}, y)}{\tilde{s}_s(y_{s-}, y)},$$

we have

$$\begin{aligned} & \mathbb{E} \left[\log Z_T^{-1} \left[\frac{\hat{\mu}^\theta}{\mu} \right] \right] \\ &= \mathbb{E} \left[- \int_0^T \int_{\mathbb{X}} \log \frac{\tilde{s}_s^\theta(y_{s-}, y)}{\tilde{s}_s(y_{s-}, y)} N[\mu](ds \times dy) + \int_0^T \int_{\mathbb{X}} \left(\frac{\tilde{s}_s^\theta(y_{s-}, y)}{\tilde{s}_s(y_{s-}, y)} - 1 \right) \mu_s(y) \nu(dy) ds \right] \\ &= \mathbb{E} \left[\int_0^T \int_{\mathbb{X}} \left(\frac{\tilde{s}_s^\theta(y_{s-}, y)}{\tilde{s}_s(y_{s-}, y)} - 1 - \log \frac{\tilde{s}_s^\theta(y_{s-}, y)}{\tilde{s}_s(y_{s-}, y)} \right) \tilde{s}_s(y_{s-}, y) \tilde{Q}_s(y_{s-}, y) \nu(dy) ds \right] \\ &= \mathbb{E} \left[\int_0^T \int_{\mathbb{X}} \left(\frac{\tilde{s}_s^\theta(y_{s-}, y)}{\tilde{s}_s(y_{s-}, y)} - \tilde{s}_s(y_{s-}, y) - \tilde{s}_s(y_{s-}, y) \log \frac{\tilde{s}_s^\theta(y_{s-}, y)}{\tilde{s}_s(y_{s-}, y)} \right) \tilde{Q}_s(y_{s-}, y) \nu(dy) ds \right], \end{aligned}$$

rearranging the terms in the above equation yields the desired result. \square

A.4 PROOFS OF STOCHASTIC INTEGRAL FORMULATIONS

Proof of Proposition 3.2. In the following, we will denote the trajectory obtained by simulating the master equation (2.7) of the forward process of the discrete diffusion model as \tilde{x}_t and the trajectory obtained by the stochastic integral (3.1) as x_t , with $x_0 = \tilde{x}_0$. We will also use the notation $\tilde{\cdot}$ to denote the quantities associated with the trajectory \tilde{x}_t . The goal is to show that x_t and \tilde{x}_t are identically distributed for any $t \in [0, T]$.

We prove this claim by induction. We assume that for any $t \in [0, t_n]$, where $n \in \mathbb{N}$ and t_n is the n -th jump time with $t_0 = 0$, the two trajectories x_t and \tilde{x}_t are identically distributed. For simplicity, we realign the two processes x_t and \tilde{x}_t at time t_n by setting $x_{t_n} = \tilde{x}_{t_n}$.

We first consider the process \tilde{x}_t generated by the discrete diffusion model (2.7). Recall the definition $\lambda_t(y) = \tilde{Q}_t(y, \tilde{x}_{t-})$, we have that

$$\int_{\mathbb{X}} \lambda_t(y) \nu(dy) = \sum_{y \in \mathbb{X}} \tilde{Q}_t(y, \tilde{x}_{t-}) = -Q_t(\tilde{x}_{t-}, \tilde{x}_{t-}) = -Q_t(\tilde{x}_{t_n}, \tilde{x}_{t_n}), \text{ for } t \in (t_n, t_{n+1}].$$

By Lemma A.11, the time interval $\Delta \tilde{t}_n = \tilde{t}_{n+1} - \tilde{t}_n$ is distributed according to (A.10), i.e.

$$\mathbb{P}(\Delta \tilde{t}_n > \tau) = \exp \left(\int_0^\tau Q_{\tau'}(\tilde{x}_{t_n}, \tilde{x}_{t_n}) d\tau' \right) = \exp \left(- \int_0^\tau \int_{\mathbb{X}} \lambda_{\tau'}(y) \nu(dy) d\tau' \right).$$

Similarly, the jump location $\tilde{x}_{t_{n+1}}$ is distributed according to (A.11), i.e.

$$\mathbb{P}(\tilde{x}_{t_{n+1}} = y) = - \frac{Q_{t_{n+1}}(y, \tilde{x}_{t_n})}{Q_{t_{n+1}}(\tilde{x}_{t_n}, \tilde{x}_{t_n})} = \frac{\lambda_{t_{n+1}}(y)}{\int_{\mathbb{X}} \lambda_{t_{n+1}}(y) \nu(dy)}.$$

Now we turn to the stochastic integral (3.1). By definition of the Poisson random measure, we have

$$\mathbb{P}(\Delta t_n > \tau) = \mathbb{P}(N[\lambda]((t_n, t_n + \tau] \times \mathbb{X}) = 0) = \exp \left(- \int_{t_n}^{t_n + \tau} \int_{\mathbb{X}} \lambda_{\tau'}(y) \nu(dy) d\tau' \right), \quad (\text{A.13})$$

and the jump location is distributed according to (A.14), *i.e.*

$$\mathbb{P}(x_{t_{n+1}} = y) = \frac{\lambda_{t_{n+1}}(y)}{\int_{\mathbb{X}} \lambda_{t_{n+1}}(y) \nu(dy)}. \quad (\text{A.14})$$

Comparing the arguments above, we conclude that the two processes x_t and \tilde{x}_t are identically distributed for any $t \in [0, t_{n+1}]$, and the induction is complete.

The proof of the equivalence between the backward process of the discrete diffusion model governed by (2.8) and the corresponding stochastic integral (3.2) can be conducted similarly, and the result follows. \square

A.4.1 τ -LEAPING

Proof of Proposition 4.1. Without loss of generality, we give the proof for $s = s_N$, and the general case can be proved similarly.

The stochastic integral (4.2) can be partitioned by the time discretization $(s_i)_{i \in [0:N]}$ into N intervals along which the evolving intensity is constant, *i.e.*

$$\begin{aligned} \hat{y}_{s_N} &= \hat{y}_0 + \int_0^s \int_{\mathbb{X}} (y - \hat{y}_{[s]^-}) N[\hat{\mu}_{[\cdot]}](ds, dy) \\ &= \hat{y}_0 + \sum_{i=1}^N \int_{s_{i-1}}^{s_i} \int_{\mathbb{X}} (y - \hat{y}_{s_{i-1}^-}) N[\hat{\mu}_{s_{i-1}}](ds, dy) \\ &= \hat{y}_0 + \sum_{i=1}^N \int_{\mathbb{X}} (y - \hat{y}_{s_{i-1}^-}) N[\hat{\mu}_{s_{i-1}}]((s_{i-1}, s_i], dy), \end{aligned}$$

which given \mathbb{X} is finite, can be further decomposed into the following sum of jumps:

$$\begin{aligned} \hat{y}_{s_N} &= \hat{y}_0 + \sum_{i=1}^N \int_{\mathbb{X}} (y - \hat{y}_{s_{i-1}^-}) N[\hat{\mu}_{s_{i-1}}]((s_{i-1}, s_i], dy) \\ &= \hat{y}_0 + \sum_{i=1}^N \sum_{y \in \mathbb{X}} (y - \hat{y}_{s_{i-1}^-}) N[\hat{\mu}_{s_{i-1}}]((s_{i-1}, s_i], \{y\}) \\ &\sim \hat{y}_0 + \sum_{i=1}^N \sum_{y \in \mathbb{X}} (y - \hat{y}_{s_{i-1}^-}) \mathcal{P}((s_i - s_{i-1}) \hat{\mu}_{s_{i-1}}(y)), \end{aligned}$$

which is exactly (4.1) in the τ -leaping algorithm (Algorithm 1). \square

Remark A.13. In Algorithm 1 and Proposition 4.1, we have implicitly assumed certain underlying additive group structure on the state space \mathbb{X} such that the difference $y - \hat{y}_{s_n}$ in (4.1) and $y - \hat{y}_{[s]^-}$ in (4.2) and their multiples are well-defined and can be summed up with states. In certain spaces with apparent ordering, e.g. $[S]^d$, this assumption can be easily satisfied by resolving violations at the boundary of the state space with clipping operations. In practice, these discrepancies only happen with a small probability given a sufficiently small stepsize, as we will show in the following proposition.

In order to quantify the issue in Remark A.13 and apply the change of measure argument in Theorem 3.3 to the τ -leaping algorithm, we provide the following alternative stochastic integral formulation of the τ -leaping algorithm.

Proposition A.14 (Alternative Stochastic Integral Formulation of τ -Leaping). *Under the assumption on the time discretization scheme and the choice of parameters in Theorem 4.7, with probability $1 - \mathcal{O}(\kappa \bar{D}^2 T)$, the τ -leaping algorithm (Algorithm 1) is equivalent to solving the following stochastic integral equation:*

$$\hat{y}_s = \hat{y}_0 + \int_0^s \int_{\mathbb{X}} (y - \hat{y}_{\tau^-}) N[\hat{\mu}_{[\cdot]}^\theta](d\tau, dy), \quad (\text{A.15})$$

where the evolving intensity $\hat{\mu}_\tau^\theta(y)$ is given by $\hat{\mu}_{[\tau]}^\theta(y) = \tilde{s}_{[\tau]}^\theta(\hat{y}_{\tau^-}, y) \tilde{Q}_{[\tau]}(\hat{y}_{\tau^-}, y)$.

Proof. For simplicity, we first consider the stochastic integral (4.2) within the time interval $(s_n, s_{n+1}]$ and assume that \hat{y} is left-continuous at time s_n which happens with probability 1.

Notice that (A.15) only differs from (4.2) in that the integrand is evaluated at each time s instead of the starting time of the interval $\lfloor s \rfloor = s_n$. Consequently, \hat{y}_{s-} coincides with $\hat{y}_{s_n-} = \hat{y}_{s_n}$ before the first jump within the interval. Therefore, both the first jump time (A.13) and the first jump location (A.14) are distributed identically for the two stochastic integral formulations, and if no further jump occurs within the interval, the two formulations are equivalent within the interval.

We bound the probability of the scenario where the two formulations differ, *i.e.* at least two jumps occur within the interval $(s_n, s_{n+1}]$. We will denote this event as A_n . By the equivalence between Algorithm 1 and the stochastic integral (4.2), we have that

$$\begin{aligned} \mathbb{P}(A_n) &= \mathbb{P} \left(\sum_{y \in \mathbb{X}} \mathcal{P}(\hat{\mu}_{s_n}^\theta(y)(s_{n+1} - s_n)) > 1 \right) \\ &= 1 - \mathbb{P} \left(\mathcal{P} \left(\sum_{y \in \mathbb{X}} \hat{\mu}_{s_n}^\theta(y)(s_{n+1} - s_n) \right) \leq 1 \right) \\ &= 1 - \exp \left(- \sum_{y \in \mathbb{X}} \hat{\mu}_{s_n}^\theta(y)(s_{n+1} - s_n) \right) \left(1 + \sum_{y \in \mathbb{X}} \hat{\mu}_{s_n}^\theta(y)(s_{n+1} - s_n) \right) \\ &\lesssim \left(\int_{s_n}^{s_{n+1}} \int_{\mathbb{X}} \hat{\mu}_{s_n}^\theta(y) \nu(dy) ds \right)^2, \end{aligned}$$

which, with a similar argument as in (C.1), is further bounded by

$$\mathbb{P}(A_n) \lesssim (1 \vee (T - s_{n+1})^{-2}) \bar{D}^2 (s_{n+1} - s_n)^2,$$

and thus

$$\begin{aligned} \mathbb{P} \left(\bigcup_{n=1}^N A_n \right) &\lesssim \sum_{n=1}^N (1 \vee (T - s_{n+1})^{-2}) \bar{D}^2 (s_{n+1} - s_n)^2 \\ &\leq \sum_{n=1}^N (1 \vee (T - s_{n+1})^{-1+\gamma-\eta}) \bar{D}^2 \kappa (s_{n+1} - s_n) \\ &\lesssim \kappa \bar{D}^2 \left(T + \int_{\delta}^1 t^{-1+\gamma-\eta} dt \right) \lesssim \kappa \bar{D}^2 T, \end{aligned}$$

where the last inequality is by following the same argument as in the proof of Proposition C.5 and the choice of the early stopping time δ in Theorem 4.7. The result follows. \square

A.4.2 UNIFORMIZATION

Theorem A.15 (Accurate Simulation by Uniformization). *The uniformization algorithm (Algorithm 2) with its stochastic integral formulation in (4.4) is equivalent to the approximate backward process with its stochastic integral formulation in (A.12).*

Proof of Proposition 4.2 and Theorem A.15. For simplicity, we only consider the stochastic integral (4.4) within the time interval $(s_b, s_{b+1}]$.

We rewrite the stochastic integral (4.4) as a sum of jumps:

$$y_{s_{b+1}} = y_{s_b} + \sum_{i=1}^N (Y_n - y_{s_{b,n}^-}) \mathbf{1}_{0 \leq \Xi_n \leq \int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^\theta(y) \nu(dy)},$$

where by Proposition A.8, $(s_{b,n})_{n \in [N]}$ are the jump times and $(Y_n, \Xi_n)_{n \in [N]}$ are the jump locations that are distributed according to

$$\mathbb{P}(Y_n = y, \Xi_n = \xi) = \frac{\hat{\mu}_{s_{b,n}}^\theta(y)}{\int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^\theta(y) \nu(dy) \int_0^{\bar{\lambda}} d\xi} = \frac{\hat{\mu}_{s_{b,n}}^\theta(y)}{\int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^\theta(y) \nu(dy) \bar{\lambda}}. \quad (\text{A.16})$$

Therefore, the n -th jump is not performed if $\int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^{\theta}(y) \nu(dy) < \Xi_n \leq \bar{\lambda}$, which is of probability

$$\begin{aligned} \mathbb{P}(\Xi_n > \hat{\mu}_{s_{b,n}}^{\theta}(y_{s_{b,n}}^-)) &= \frac{\int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^{\theta}(y) \nu(dy) \bar{\lambda} - \int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^{\theta}(y) \nu(dy)}{\bar{\lambda}} \\ &= 1 - \frac{\int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^{\theta}(y) \nu(dy)}{\bar{\lambda}}, \end{aligned}$$

and is to the state y with probability

$$\begin{aligned} \mathbb{P}\left(Y_n = y, \Xi_n \leq \int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^{\theta}(y) \nu(dy)\right) \\ = \frac{\hat{\mu}_{s_{b,n}}^{\theta}(y)}{\int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^{\theta}(y) \nu(dy)} \frac{\int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^{\theta}(y) \nu(dy)}{\bar{\lambda}} = \frac{\hat{\mu}_{s_{b,n}}^{\theta}(y)}{\bar{\lambda}}, \end{aligned}$$

which coincides with (4.3) in the uniformization algorithm (Algorithm 2).

By conditioning on the occurrence of each jump, *i.e.* $0 \leq \Xi_n \leq \int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^{\theta}(y) \nu(dy)$, with slight abuse of notation, we have that

$$\mathbb{P}\left(Y_n = y \mid 0 \leq \Xi_n \leq \int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^{\theta}(y) \nu(dy)\right) = \frac{\hat{\mu}_{s_{b,n}}^{\theta}(y)}{\int_{\mathbb{X}} \hat{\mu}_{s_{b,n}}^{\theta}(y) \nu(dy)},$$

which again by Proposition A.8 implies that y_s also satisfies the stochastic integral (A.12) corresponding to the approximate backward process, and vice versa, and the result follows. \square

B RESULTS FOR CONTINUOUS-TIME MARKOV CHAIN

In this section, we will provide some results for the continuous-time Markov chain (CTMC), including the mixing time, the spectral gap, the modified log-Sobolev constant, *etc.*

Definition B.1 (Graph Corresponding to Rate Matrix). *We denote $\mathcal{G}(\mathbf{Q})$ as the graph corresponding to the rate matrix \mathbf{Q} , i.e.*

$$\mathcal{G}(\mathbf{Q}) = (\mathbb{X}, E(\mathcal{G}(\mathbf{Q})), \mathbf{Q}), \text{ where } E(\mathcal{G}(\mathbf{Q})) = \{(x, y) \in \mathbb{X} \times \mathbb{X} \mid x \neq y, \mathbf{Q}(x, y) > 0\},$$

and the weight of the directed edge $(x, y) \in E(\mathcal{G}(\mathbf{Q}))$ is $\mathbf{Q}(x, y)$.

We will assume that the continuous-time Markov chain is irreducible and reversible on the state space \mathbb{X} , and the corresponding stationary distribution is π .

B.1 SPECTRAL GAP

Definition B.2 (Spectral Gap). *Let $\mathbf{L} = -\mathbf{Q}$ be the graph Laplacian matrix with $\mathbf{D} = \text{diag } \mathbf{L}$, corresponding to the graph $\mathcal{G}(\mathbf{Q})$. with*

$$0 = \lambda_1(\mathbf{L}) < \lambda_2(\mathbf{L}) \leq \dots \leq \lambda_{|\mathbb{X}|}(\mathbf{L}) \leq 2 \max_{x \in \mathbb{X}} D(x, x) = 2\bar{D},$$

the spectral gap $\lambda(\mathbf{Q})$ of the rate matrix \mathbf{Q} is defined as the second smallest eigenvalue of the graph Laplacian \mathbf{L} , i.e. $\lambda(\mathbf{Q}) = \lambda_2(\mathbf{L})$.

Remark B.3 (Asymptotic Behavior of the score function s_t). *Assume \mathbf{Q} is symmetric with the following orthogonal eigendecomposition:*

$$\mathbf{Q} = -\mathbf{U}\mathbf{\Lambda}\mathbf{U}^{\top},$$

where $\mathbf{U} = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_{|\mathbb{X}|})$ is an orthogonal matrix, and the distribution \mathbf{p}_t has the following decomposition w.r.t. the eigenvectors of the graph Laplacian \mathbf{L} :

$$\mathbf{p}_t = \sum_{i=1}^{|\mathbb{X}|} \alpha_t(i) \mathbf{u}_i = \mathbf{U} \boldsymbol{\alpha}(t),$$

then the solution to the master equation (2.7) is given by

$$\mathbf{p}_t = \exp(t\mathbf{Q})\mathbf{p}_0 = \mathbf{U} \exp(-t\mathbf{\Lambda})\mathbf{U}^\top \mathbf{p}_0 = \mathbf{U} \exp(-t\mathbf{\Lambda})\boldsymbol{\alpha}_0 = \sum_{j=1}^{|\mathbb{X}|} \mathbf{u}_j \exp(-t\lambda_j) \alpha_0(j),$$

i.e. $\boldsymbol{\alpha}_t = \exp(-t\mathbf{\Lambda})\boldsymbol{\alpha}_0$ and thus for any $i \in [|\mathbb{X}|]$,

$$p_t(i) - p_0(i) = \sum_{j=1}^{|\mathbb{X}|} u_j(i) (-1 + \exp(-t\lambda_j)) \alpha_0(j) = - \sum_{j>1} u_j(i) \alpha_0(j) \lambda_j \mathcal{O}(t).$$

Therefore, we have

$$s_t(x, y) = \frac{p_t(y)}{p_t(x)} = \frac{p_0(y) - \sum_{j>1} u_j(y) \alpha_0(j) \lambda_j \mathcal{O}(t)}{p_0(x) - \sum_{j>1} u_j(x) \alpha_0(j) \lambda_j \mathcal{O}(t)} \lesssim 1 \vee (Ft)^{-1},$$

given that the following condition is satisfied

$$F = \min_{x \in \mathbb{X}} \left| \sum_{j>1} u_j(x) \alpha_0(j) \lambda_j \right| > 0,$$

which only depends on the initial distribution \mathbf{p}_0 and the rate matrix \mathbf{Q} .

Epecially, the bound $s_t(x, y) \lesssim 1$ for any $x \in \mathbb{X}$ s.t. $p_0(x) > 0$ and $s_t(x, y) \lesssim t^{-1}$ for those s.t. $p_0(x) = 0$.

Definition B.4 (Conductance). The conductance $\phi(\mathcal{G})$ of a graph \mathcal{G} is defined as

$$\phi(\mathcal{G}) = \min_{S \subset \mathbb{X}} \frac{\sum_{x \in S, y \notin S} Q(x, y)}{\min \left\{ \sum_{x \in S} D(x, x), \sum_{y \notin S} D(y, y) \right\}}.$$

Theorem B.5 (Cheeger's Inequality). Denote the normalized graph Laplacian matrix by $\tilde{\mathbf{L}} = \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}$ with eigenvalues

$$0 \leq \lambda_1(\tilde{\mathbf{L}}) \leq \lambda_2(\tilde{\mathbf{L}}) \leq \dots \leq \lambda_{|\mathbb{X}|}(\tilde{\mathbf{L}}) \leq 2,$$

then the conductance of the graph $\mathcal{G}(\mathbf{Q})$ can be bounded by

$$\frac{1}{2} \lambda_2(\tilde{\mathbf{L}}) \leq \phi(\mathcal{G}(\mathbf{Q})) \leq \sqrt{2 \lambda_2(\tilde{\mathbf{L}})}.$$

B.2 LOG-SOBOLEV INEQUALITIES

Definition B.6 (Modified Log-Sobolev Constant (Bobkov & Tetali, 2006)). For any function $f, g : \mathbb{X} \rightarrow \mathbb{R}$, we denote the entropy functional $\text{Ent}_\pi(f)$ of f as

$$\text{Ent}_\pi(f) := \mathbb{E}_\pi[f \log f] - \mathbb{E}_\pi[f] \log \mathbb{E}_\pi[f],$$

and the Dirichlet form $\mathcal{E}_\pi(f, g)$ as

$$\mathcal{E}_\pi(f, g) = \mathbb{E}_\pi[f \mathbf{L}^T g] := \sum_{y \in \mathbb{X}} f(y) (\mathbf{L}^T g)(y) \pi(y) = \sum_{x, y \in \mathbb{X}} f(y) L(x, y) g(x) \pi(y),$$

where the Laplacian $\mathbf{L} = \mathbf{Q}$. Then the modified log-Sobolev constant of the rate matrix \mathbf{Q} is defined as

$$\rho(\mathbf{Q}) := \inf \left\{ \frac{\mathcal{E}_\pi(f, \log f)}{\text{Ent}_\pi(f)} \mid f : \mathbb{X} \rightarrow \mathbb{R}, \text{Ent}_\pi(f) > 0 \right\}.$$

Theorem B.7 (Theorem 2.4, (Bobkov & Tetali, 2006)). For any initial distribution \mathbf{p}_0 , we have for any $t \geq 0$,

$$D_{\text{KL}}(\mathbf{p}_t \| \boldsymbol{\pi}) \leq D_{\text{KL}}(\mathbf{p}_0 \| \boldsymbol{\pi}) \exp(-\rho(\mathbf{Q})t),$$

i.e. the KL divergence converges exponentially fast with rate $\rho(\mathbf{Q})$.

Proof. Noticing that $\text{Ent}_\pi(\frac{p_t}{\pi}) = D_{\text{KL}}(\mathbf{p}_t \parallel \pi)$, we differentiate $D_{\text{KL}}(\mathbf{p}_t \parallel \pi)$ with respect to t to obtain that

$$\begin{aligned} \frac{d}{dt} D_{\text{KL}}(\mathbf{p}_t \parallel \pi) &= \frac{d}{dt} \sum_{x \in \mathbb{X}} \frac{p_t(x)}{\pi(x)} \log \left(\frac{p_t(x)}{\pi(x)} \right) \pi(x) = \sum_{x \in \mathbb{X}} \left(\log \frac{p_t(x)}{\pi(x)} + 1 \right) \pi(x) \frac{d}{dt} \frac{p_t(x)}{\pi(x)} \\ &= \sum_{x \in \mathbb{X}} \left(\log \frac{p_t(x)}{\pi(x)} + 1 \right) \frac{d}{dt} p_t(x) = - \sum_{x, y \in \mathbb{X}} \log \frac{p_t(x)}{\pi(x)} L(x, y) p_t(y) \\ &= - \sum_{y \in \mathbb{X}} \frac{p_t(y)}{\pi(y)} \left(\sum_{x \in \mathbb{X}} L(x, y) \log \frac{p_t(x)}{\pi(x)} \right) \pi(y) \\ &= - \mathcal{E}_\pi \left(\frac{p_t}{\pi}, \log \frac{p_t}{\pi} \right) \leq -\rho(\mathbf{Q}) D_{\text{KL}}(\mathbf{p}_t \parallel \pi), \end{aligned} \quad (\text{B.1})$$

and the result follows by applying Grönwall's inequality to both sides above. \square

Then, the following proposition connects the modified log-Sobolev constant with the spectral gap.

Proposition B.8 ((Bobkov & Tetali, 2006, Proposition 3.5)). *The modified log-Sobolev constant $\rho(\mathbf{Q})$ of the rate matrix \mathbf{Q} is bounded by the spectral gap $\lambda(\mathbf{Q})$, i.e. $\rho(\mathbf{Q}) \leq \lambda(\mathbf{Q})$.*

Proof. Below we provide a sketch of the informal proof of the proposition above for the sake of completeness. Let $f : \mathbb{X} \rightarrow \mathbb{R}$ be an arbitrary function and $\zeta > 0$ be any positive number. From the definition of the modified log-Sobolev constant, we have

$$\mathcal{E}_\pi(e^{\zeta f}, \zeta f) \geq \rho(\mathbf{Q}) \text{Ent}_\pi(e^{\zeta f}). \quad (\text{B.2})$$

Under the limit $\zeta \rightarrow 0^+$, we may apply Taylor expansion to the two terms on the LHS and RHS, which implies

$$\begin{aligned} \mathcal{E}_\pi(e^{\zeta f}, \zeta f) &= \mathcal{E}_\pi(1 + \zeta f + O(\zeta^2), \zeta f) \\ &= \zeta \mathcal{E}_\pi(1, f) + \zeta^2 \mathcal{E}_\pi(f, f) + O(\zeta^3) = \zeta^2 \mathcal{E}_\pi(f, f) + O(\zeta^3) \\ \text{Ent}_\pi(e^{\zeta f}) &= \mathbb{E}_\pi[e^{\zeta f} \zeta f] - \mathbb{E}_\pi[e^{\zeta f}] \log \mathbb{E}_\pi[e^{\zeta f}] \\ &= \mathbb{E}_\pi[(1 + \zeta f + O(\zeta^2)) \zeta f] - \mathbb{E}_\pi[1 + \zeta f + O(\zeta^2)] \log \mathbb{E}_\pi[e^{\zeta f}] \\ &= \zeta \mathbb{E}_\pi[f] + \zeta^2 \mathbb{E}_\pi[f^2] + O(\zeta^3) \\ &\quad - (1 + \zeta \mathbb{E}_\pi[f] + O(\zeta^2)) \log(1 + \zeta \mathbb{E}_\pi[f] + O(\zeta^2)) \\ &= \zeta \mathbb{E}_\pi[f] + \zeta^2 \mathbb{E}_\pi[f^2] + O(\zeta^3) - (1 + \zeta \mathbb{E}_\pi[f] + O(\zeta^2)) (\zeta \mathbb{E}_\pi[f] + O(\zeta^2)) \\ &= \zeta^2 (\mathbb{E}_\pi[f^2] - \mathbb{E}_\pi[f]^2) + O(\zeta^3) \end{aligned} \quad (\text{B.3})$$

Substituting (B.3) into (B.2) and taking the limit $\zeta \rightarrow 0^+$ then yield the following inequality

$$\rho(\mathbf{Q}) \leq \frac{\mathcal{E}_\pi(f, f)}{\mathbb{E}_\pi[f^2] - \mathbb{E}_\pi[f]^2}$$

for any non-constant function $f : \mathbb{X} \rightarrow \mathbb{R}$. Taking infimum on both sides above with respect to all f then indicates

$$\rho(\mathbf{Q}) \leq \inf_f \frac{\mathcal{E}_\pi(f, f)}{\mathbb{E}_\pi[f^2] - \mathbb{E}_\pi[f]^2} \leq \inf_{f: \mathbb{E}_\pi[f]=0} \frac{\mathcal{E}_\pi(f, f)}{\mathbb{E}_\pi[f^2]} = \lambda_2(\mathbf{L}) = \lambda(\mathbf{Q}) \quad (\text{B.4})$$

where the last two equalities above follows from the definition of spectral gap, as desired. \square

In general, the lower bound of the modified log-Sobolev constant $\rho(\mathbf{Q})$ and the spectral gap $\lambda(\mathbf{Q})$ depends on the connectivity and other specific structures of the graph $\mathcal{G}(\mathbf{Q})$, and the related research is still an active area on a graph-by-graph basis (Bobkov & Tetali, 2006).

The properties of the spectral gap $\lambda(\mathbf{Q})$ are better known in the literature, as it is closely related to the conductance of the graph $\mathcal{G}(\mathbf{Q})$ via Cheeger's inequality (Theorem B.5), and thus when $\mathbf{D} = \mathbf{D}\mathbf{I}$, the spectral gap $\lambda(\mathbf{Q})$ satisfies

$$\frac{1}{2D} \lambda(\mathbf{Q}) = \frac{1}{2} \lambda_2(\mathbf{L}) \leq \phi(\mathcal{G}(\mathbf{Q})) \leq \sqrt{2\lambda_2(\mathbf{L})} = \sqrt{\frac{2\lambda(\mathbf{Q})}{D}}.$$

However, as shown in Proposition B.8, the lower bound on the modified log-Sobolev constant $\rho(\mathbf{Q})$ is hard to obtain, as the KL divergence, the exponential convergence of which is controlled by $\rho(\mathbf{Q})$, is stronger than the total variation distance, the exponential convergence of which is controlled by $\lambda(\mathbf{Q})$, via Pinsker's inequality. The following theorem gives a rough lower bound on d -regular graphs.

Theorem B.9 (Bobkov & Tetali, 2006, Proposition 5.4). *Suppose \mathcal{G} is a d -regular graph on \mathbb{X} with unit weights and \mathbf{Q} is the corresponding rate matrix such that $\mathcal{G}(\mathbf{Q}) = \mathcal{G}$, then the modified log-Sobolev constant $\rho(\mathbf{Q})$ of the rate matrix \mathbf{Q} satisfies*

$$\frac{\lambda(\mathbf{Q})}{\log |\mathbb{X}|} \leq \rho(\mathbf{Q}) \leq \frac{8d \log |\mathbb{X}|}{\text{diam}(\mathcal{G})^2},$$

where $\text{diam}(\mathcal{G})$ is the diameter of the graph \mathcal{G} .

For some specific graphs, the modified log-Sobolev constant $\rho(\mathbf{Q})$ and the spectral gap $\lambda(\mathbf{Q})$ can be explicitly calculated, such as the following examples:

Example B.10 (Hypercube (Gross, 1975)). *Let $\mathbb{X} = \{0, 1\}^d$ and \mathbf{Q} be the rate matrix for which the graph $\mathcal{G}(\mathbf{Q})$ is a hypercube, and for any two states $x, y \in \mathbb{X}$, the rate $Q(x, y) = 1$ if x and y differ in exactly one coordinate. Then the modified log-Sobolev constant $\rho(\mathbf{Q})$ and the spectral gap $\lambda(\mathbf{Q})$ are given by*

$$\rho(\mathbf{Q}) = \lambda(\mathbf{Q}) = 4,$$

which is dimensionless.

Example B.11 (Asymmetric Hypercube (Diaconis & Saloff-Coste, 1996)). *Let $\mathbb{X} = \{0, 1\}^d$ and \mathbf{Q} be the rate matrix for which the graph $\mathcal{G}(\mathbf{Q})$ is a hypercube, and for any two states $x, y \in \mathbb{X}$, the rate $Q(x, y) = p$ if x and y differ in exactly one coordinate and x is the state with 0 in that coordinate, and $Q(x, y) = q = 1 - p$ if with 1 in that coordinate. Then the modified log-Sobolev constant $\rho(\mathbf{Q})$ and the spectral gap $\lambda(\mathbf{Q})$ are given by*

$$\rho(\mathbf{Q}) = \frac{2(p - q)}{pq(\log p - \log q)}, \text{ and } \lambda(\mathbf{Q}) = \frac{1}{pq}.$$

Further results on log-Sobolev inequalities related to finite-state Markov chains are beyond the scope of this paper, and we refer the readers to (Stroock, 1993; Saloff-Coste, 1997; Bobkov & Ledoux, 1998; Lee & Yau, 1998; Goel, 2004; Ledoux, 2006) for more detail.

B.3 MIXING TIME

Definition B.12 (Mixing Time). *We define the mixing time $t_{\text{mix}}(\epsilon)$ of the continuous-time Markov chain with rate matrix \mathbf{Q} as the smallest time t such that starting from any initial distribution \mathbf{p}_0 , the KL divergence $D_{\text{KL}}(\mathbf{p}_t \| \boldsymbol{\pi})$ is less than ϵ , i.e.*

$$t_{\text{mix}}(\epsilon) = \inf \left\{ t \in \mathbb{R}_+ \mid D_{\text{KL}}(\mathbf{p}_t \| \boldsymbol{\pi}) = D_{\text{KL}}(e^{-t\mathbf{Q}}\mathbf{p}_0 \| \boldsymbol{\pi}) \leq \epsilon \right\}.$$

Similarly, we define the mixing time $t_{\text{mix,TV}}(\epsilon)$ as the smallest time t such that starting from any initial distribution \mathbf{p}_0 , the total variation distance $\text{TV}(\mathbf{p}_t, \boldsymbol{\pi})$ is less than ϵ , i.e.

$$t_{\text{mix,TV}}(\epsilon) = \inf \left\{ t \in \mathbb{R}_+ \mid \text{TV}(\mathbf{p}_t, \boldsymbol{\pi}) = \text{TV}(e^{-t\mathbf{Q}}\mathbf{p}_0, \boldsymbol{\pi}) \leq \epsilon \right\}.$$

With a slight abuse of notation, we will also denote the e^{-1} -mixing time as $t_{\text{mix}} = t_{\text{mix,KL}}(e^{-1})$ and $t_{\text{mix,TV}} = t_{\text{mix,TV}}(e^{-1})$.

Proposition B.13. *The mixing time $t_{\text{mix}}(\epsilon)$ of the continuous-time Markov chain with rate matrix \mathbf{Q} is bounded by the modified log-Sobolev constant $\rho(\mathbf{Q})$, i.e.*

$$t_{\text{mix}}(\epsilon) \lesssim \rho(\mathbf{Q})^{-1} (\log \epsilon^{-1} + \log \log \pi_*^{-1}),$$

And the mixing time $t_{\text{mix,TV}}(\epsilon)$ is bounded by the spectral gap $\lambda(\mathbf{Q})$, i.e.

$$t_{\text{mix,TV}}(\epsilon) \lesssim \lambda(\mathbf{Q})^{-1} (\log \epsilon^{-1} + \log \log \pi_*^{-1}).$$

Proof. Define $\pi_* = \min_{x \in \mathbb{X}} \pi(x)$, we first bound $D_{\text{KL}}(\mathbf{p}_0 \| \boldsymbol{\pi})$ as follows:

$$D_{\text{KL}}(\mathbf{p}_0 \| \boldsymbol{\pi}) \leq \sum_{x \in \mathbb{X}} p_0(x) \log \frac{p_0(x)}{\pi(x)} \leq \log \pi_*^{-1}, \quad (\text{B.5})$$

and thus by Theorem B.7, we have

$$D_{\text{KL}}(\mathbf{p}_t \| \boldsymbol{\pi}) \leq D_{\text{KL}}(\mathbf{p}_0 \| \boldsymbol{\pi}) \exp(-\rho(\mathbf{Q})t) \leq \log \pi_*^{-1} \exp(-\rho(\mathbf{Q})t). \quad (\text{B.6})$$

Therefore, by setting the right-hand side of (B.6) to be ϵ , we have the desired result for the mixing time

$$t_{\text{mix}}(\epsilon) \leq \frac{1}{\rho(\mathbf{Q})} (\log \epsilon^{-1} + \log \log \pi_*^{-1}).$$

For the mixing time $t_{\text{mix,TV}}(\epsilon)$, we use the Pinsker's inequality to obtain:

$$\text{TV}(\mathbf{p}_t, \boldsymbol{\pi}) \leq \sqrt{\frac{1}{2} D_{\text{KL}}(\mathbf{p}_t \| \boldsymbol{\pi})} \leq \sqrt{\frac{1}{2} \log \pi_*^{-1} \exp(-\rho(\mathbf{Q})t)},$$

and therefore,

$$t_{\text{mix,TV}}(\epsilon) \lesssim \frac{1}{\rho(\mathbf{Q})} \log \left(\epsilon^{-1} \sqrt{\log \pi_*^{-1}} \right) \lesssim \frac{1}{\rho(\mathbf{Q})} (\log \epsilon^{-1} + \log \log \pi_*^{-1}).$$

□

Corollary B.14. *The e^{-1} -mixing time t_{mix} and $t_{\text{mix,TV}}$ of the continuous-time Markov chain with rate matrix \mathbf{Q} satisfy*

$$t_{\text{mix}} \lesssim \rho(\mathbf{Q})^{-1} \log \log \pi_*^{-1}, \quad \text{and} \quad t_{\text{mix,TV}} \lesssim \rho(\mathbf{Q})^{-1} \log \log \pi_*^{-1},$$

and thus $t_{\text{mix}} \gtrsim t_{\text{mix,TV}}$.

C PROOFS OF ERROR ANALYSIS IN SECTION 4.3

In this section, we provide the proof of the primary results in the main text and their sketches.

C.1 PROOF SKETCH

A general pipeline of the proofs of Theorem 4.7 and 4.9 is to decompose the KL divergence between the target distribution and the distribution of generated sample as a summation of the truncation, approximation, and discretization errors, and then bound each term separately, echoing that for the continuous case in Section 2.3.

- **Truncation Error:** As in the continuous case, the truncation error is caused by the finite time horizon of the forward process (2.7) and is thus bounded by the convergence rate of the constructed forward process. In our work, we use the modified log-Sobolev constant $\rho(\mathbf{Q})$ to bound the truncation error (cf. Definition B.6). As discussed in Appendix B, a lower bound on the modified log-Sobolev constant $\rho(\mathbf{Q})$ guarantees the exponential convergence of the forward process (cf. Theorem B.7), which then implies the upper bound on the truncation error presented in Proposition C.1.
- **Discretization Error:** In the τ -leaping algorithm, the discretization error is caused by the approximation of the backward process with a piecewise constant function, and should decrease as the time step κ decreases. We first analyze the discretization error within one step (cf. Proposition C.4) and then provide the overall discretization error in Proposition C.5, taking different discretization scheme into consideration and deriving early stopping criteria. In the uniformization algorithm, due to its exactness, the discretization error is zero.
- **Approximation Error:** The approximation error is caused by the estimation error of the score function, and is bounded by certain error assumptions (cf. Assumption 4.6 for τ -leaping and Assumption 4.6' for uniformization) on the score function estimation in the training process.

Then in Appendix C.4, we first invoke the change of measure theorem 3.3 and obtain KL divergence bounds for both algorithms (Theorem C.6 and C.7), and then provide the overall error bounds thereafter.

C.2 TRUNCATION ERROR

Theorem C.1 (Truncation Error). *The forward process (2.7) converges to the uniform distribution $\mathbf{p}_\infty = \mathbf{1}/|\mathbb{X}|$ exponentially fast in terms of the KL divergence, i.e.*

$$D_{\text{KL}}(\mathbf{p}_t \parallel \mathbf{p}_\infty) = D_{\text{KL}}\left(\mathbf{p}_t \parallel \frac{\mathbf{1}}{|\mathbb{X}|}\right) \lesssim e^{-\rho t} \log |\mathbb{X}|,$$

where $|\mathbb{X}|$ is the size of the state space, and t_{mix} is the mixing time of the continuous-time Markov chain corresponding to the rate matrix \mathbf{Q} defined in Definition B.12.

Proof. Since \mathbf{Q} is symmetric, we have the stationary distribution $\pi = \mathbf{1}/|\mathbb{X}|$ and thus $D_{\text{KL}}(\mathbf{p}_t \parallel \mathbf{p}_\infty) = D_{\text{KL}}(\mathbf{p}_t \parallel \pi)$, and $\pi_* = 1/|\mathbb{X}|$. By Assumption 4.3 and Theorem B.7, we have

$$D_{\text{KL}}(\mathbf{p}_t \parallel \pi) \leq e^{-\rho(\mathbf{Q})t} D_{\text{KL}}(\mathbf{p}_0 \parallel \pi) \leq e^{-\rho t} \log \pi_*^{-1} \leq e^{-\rho t} \log |\mathbb{X}|,$$

where the last inequality is by (B.5). \square

C.3 DISCRETIZATION ERROR

Denote the shorthand notation $G(x; y) = x(\log x - \log y) - x$; it is easy to check that $G'(x; y) = \log x - \log y$.

Proposition C.2. *For any $y \in \mathbb{X}$, we have*

$$|\partial_\sigma G(\mu_\sigma(y); \hat{\mu}_{[s_n]}^\theta(y))| \lesssim (\log C + \log(1 \vee (T - \sigma)) + \log M) \mu_\sigma(y) \bar{D} (1 \vee (T - \sigma)^{-2}).$$

Proof. By the chain rule, we have

$$\partial_\sigma G(\mu_\sigma(y); \hat{\mu}_{[s_n]}^\theta(y)) = G'(\mu_\sigma(y); \hat{\mu}_{[s_n]}^\theta(y)) \partial_\sigma \mu_\sigma(y) = \left(\log \mu_\sigma(y) - \log \hat{\mu}_{[s_n]}^\theta(y) \right) \partial_\sigma \mu_\sigma(y).$$

We first compute $\partial_\sigma \mu_\sigma(y)$ as

$$\begin{aligned} \partial_\sigma \mu_\sigma(y) &= \tilde{Q}(\tilde{x}_{\sigma-}, y) \partial_\sigma \tilde{s}_\sigma(\tilde{x}_{\sigma-}, y) = \tilde{Q}(\tilde{x}_{\sigma-}, y) \partial_\sigma \left(\frac{\tilde{p}_\sigma(y)}{\tilde{p}_\sigma(x_{\sigma-})} \right) \\ &= \tilde{Q}(\tilde{x}_{\sigma-}, y) \left(\frac{1}{\tilde{p}_\sigma(x_{\sigma-})} \partial_\sigma \tilde{p}_\sigma(y) - \frac{\tilde{p}_\sigma(y)}{\tilde{p}_\sigma(x_{\sigma-})^2} \partial_\sigma \tilde{p}_\sigma(x_{\sigma-}) \right) \\ &= \tilde{Q}(\tilde{x}_{\sigma-}, y) \left(-\frac{\tilde{p}_\sigma(y)}{\tilde{p}_\sigma(x_{\sigma-})} \sum_{y' \in \mathbb{X}} \frac{\tilde{p}_\sigma(y')}{\tilde{p}_\sigma(y)} Q(y, y') + \frac{\tilde{p}_\sigma(y)}{\tilde{p}_\sigma(x_{\sigma-})} \sum_{y' \in \mathbb{X}} \frac{\tilde{p}_\sigma(y')}{\tilde{p}_\sigma(x_{\sigma-})} Q(x_{\sigma-}, y') \right) \\ &= \mu_\sigma(y) \left(-\sum_{y' \in \mathbb{X}} \tilde{s}_\sigma(y, y') Q(y, y') + \sum_{y' \in \mathbb{X}} \tilde{s}_\sigma(x_{\sigma-}, y') Q(x_{\sigma-}, y') \right), \end{aligned}$$

by which we have

$$|\partial_\sigma \mu_\sigma(y)| \lesssim \mu_\sigma(y) \left(\sum_{y' \in \mathbb{X}} (1 \vee (T - \sigma)^{-2}) |Q(x_{\sigma-}, y')| \right) \lesssim \mu_\sigma(y) \bar{D} (1 \vee (T - \sigma)^{-2}),$$

and thus

$$\begin{aligned} \left| \partial_\sigma G(\mu_\sigma(y); \hat{\mu}_{[s_n]}^\theta(y)) \right| &\leq \left| \log \mu_\sigma(y) - \log \hat{\mu}_{[s_n]}^\theta(y) \right| |\partial_\sigma \mu_\sigma(y)| \\ &\leq \left(|\log \tilde{Q}(\tilde{x}_{\sigma-}, y)| + |\log \tilde{s}_\sigma(\tilde{x}_{\sigma-}, y)| + |\log \tilde{Q}(\tilde{x}_{\sigma-}, y)| + |\log \tilde{s}_{s_n}^\theta(\tilde{x}_{\sigma-}, y)| \right) |\partial_\sigma \mu_\sigma(y)| \\ &\lesssim \mu_\sigma(y) (\log C + \log(1 \vee (T - \sigma)^{-1}) + \log M) \bar{D} (1 \vee (T - \sigma)^{-2}). \end{aligned}$$

\square

Proposition C.3. For any $0 < s < t \leq T$, we have

$$\int_s^t \int_{\mathbb{X}} \mu_\sigma(y') \nu(dy') d\sigma \lesssim (1 \vee (T-t)^{-1}) \overline{D}(t-s).$$

Proof.

$$\begin{aligned} \int_s^t \int_{\mathbb{X}} \mu_\sigma(y') \nu(dy') d\sigma &= \int_s^t \int_{\mathbb{X}} \tilde{s}_\sigma(\tilde{x}_{\sigma-}, y') \tilde{Q}(\tilde{x}_{\sigma-}, y') \nu(dy') d\sigma \\ &\lesssim (1 \vee (T-t)^{-1}) \int_s^t \int_{\mathbb{X}} \tilde{Q}(y, \tilde{x}_{\sigma-}) \nu(dy') d\sigma \\ &\lesssim (1 \vee (T-t)^{-1}) \int_s^t |Q(\tilde{x}_{\sigma-}, \tilde{x}_{\sigma-})| d\sigma \lesssim (1 \vee (T-t)^{-1}) \overline{D}(t-s). \end{aligned} \tag{C.1}$$

□

Proposition C.4. For any $y \in \mathbb{X}$, we have

$$\begin{aligned} &\mathbb{E} \left[\left| G(\mu_s; \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y)) - G(\mu_{s_n}; \hat{\mu}_{\lfloor s \rfloor}^\theta(y)) \right| \right] \\ &\lesssim (\log C + \log(1 \vee (T-s_{n+1})^{-1}) + \log M) \mu_\sigma(y) \overline{D}(s-s_n) (1 \vee (T-s_{n+1})^{-1-\gamma}). \end{aligned}$$

Proof. Applying Theorem A.10 to the backward process (3.2), we have

$$\begin{aligned} G(\mu_s(y); \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y)) &= G(\mu_{s_n}(y); \hat{\mu}_{\lfloor s \rfloor}^\theta(y)) + \int_{s_n}^s \partial_\sigma G(\mu_\sigma(y); \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y)) d\sigma \\ &\quad + \int_{s_n}^s \int_{\mathbb{X}} \left(G(\mu_{\sigma+}(y); \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y)) - G(\mu_\sigma(y); \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y)) \right) N[\mu](d\sigma, dy'), \end{aligned}$$

where we adopt the notation $\mu_{\sigma+}(y)$ as the right limit of the càglàd process $\mu_\sigma(y)$, i.e. $\mu_{\sigma+}(y) = \tilde{s}_\sigma(\tilde{x}_\sigma, y) \tilde{Q}(\tilde{x}_\sigma, y)$.

For the first term, we have by Proposition C.2 that

$$\begin{aligned} &\left| \int_{s_n}^s \partial_\sigma G(\mu_\sigma(y); \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y)) d\sigma \right| \lesssim \int_{s_n}^s \left| \partial_\sigma G(\mu_\sigma(y); \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y)) \right| d\sigma \\ &\lesssim \int_{s_n}^s (\log C + \log(1 \vee (T-\sigma)^{-1}) + \log M) \mu_\sigma(y) \overline{D}(1 \vee (T-\sigma)^{-2}) d\sigma \\ &\lesssim (\log C + \log(1 \vee (T-s_{n+1})^{-1}) + \log M) \mu_\sigma(y) \overline{D}(s-s_n) (1 \vee (T-s_{n+1})^{-1}). \end{aligned}$$

For the second term, we have

$$\begin{aligned} &\left| G(\mu_{\sigma+}(y); \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y)) - G(\mu_\sigma(y); \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y)) \right| \\ &= \left| G'(\xi; \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y)) (\mu_{\sigma+}(y) - \mu_\sigma(y)) \right| = \left| (\log \xi - \log \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y)) (\mu_{\sigma+}(y) - \mu_\sigma(y)) \right| \\ &\leq \left(|\log \mu_{\sigma+}(y)| + |\log \mu_\sigma(y)| + \left| \log \hat{\mu}_{\lfloor s_n \rfloor}^\theta(y) \right| \right) \left| \frac{\mu_{\sigma+}(y)}{\mu_\sigma(y)} - 1 \right| \mu_\sigma(y) \\ &\lesssim (\log C + \log(1 \vee (T-\sigma)^{-1}) + \log M) (1 \vee (T-\sigma)^{-\gamma}) \mu_\sigma(y), \end{aligned}$$

where the first equality follows from the mean value theorem and the last inequality is by Assumption 4.5.

Therefore,

$$\begin{aligned}
& \mathbb{E} \left[\left| G(\mu_s; \hat{\mu}_{[s_n]}^\theta(y)) - G(\mu_{s_n}; \hat{\mu}_{[s]}^\theta(y)) \right| \right] \\
& \leq \mathbb{E} \left[\left| \int_{s_n}^s \partial_\sigma G(\mu_\sigma(y); \hat{\mu}_{[s_n]}^\theta(y)) d\sigma \right| \right] \\
& + \mathbb{E} \left[\int_{s_n}^s \int_{\mathbb{X}} \left| G(\mu_{\sigma^+}(y); \hat{\mu}_{[s_n]}^\theta(y)) - G(\mu_\sigma(y); \hat{\mu}_{[s_n]}^\theta(y)) \right| N[\mu](d\sigma, dy') \right] \\
& \lesssim (\log C + \log(1 \vee (T - s_{n+1})^{-1}) + \log M) \mu_\sigma(y) \bar{D}(s - s_n) (1 \vee (T - s_{n+1})^{-1}) \\
& + \int_{s_n}^s \int_{\mathbb{X}} (\log C + \log(1 \vee (T - \sigma)^{-1}) + \log M) \mu_\sigma(y) (1 \vee (T - \sigma)^{-\gamma}) \mu_\sigma(y') \nu(dy') d\sigma \\
& \lesssim (\log C + \log(1 \vee (T - s_{n+1})^{-1}) + \log M) \mu_\sigma(y) \bar{D}(s - s_n) (1 \vee (T - s_{n+1})^{-1}) \\
& + (\log C + \log(1 \vee (T - s_{n+1})^{-1}) + \log M) \mu_\sigma(y) (1 \vee (T - s_{n+1})^{-1-\gamma}) (s - s_n) \bar{D} \\
& \lesssim (\log C + \log(1 \vee (T - s_{n+1})^{-1}) + \log M) \mu_\sigma(y) \bar{D}(s - s_n) (1 \vee (T - s_{n+1})^{-1-\gamma}),
\end{aligned}$$

where the second to last inequality is by Proposition C.3. \square

Proposition C.5 (Discretization Error). *The following bound holds*

$$\begin{aligned}
& \int_0^{T-\delta} \int_{\mathbb{X}} \left| G(\mu_s(y); \hat{\mu}_{[s_n]}^\theta(y)) - G(\mu_{s_n}(y); \hat{\mu}_{[s]}^\theta(y)) \right| \nu(dy) ds \\
& \lesssim \begin{cases} \bar{D}^2 \kappa T, & \gamma < 1, \\ \bar{D}^2 \kappa (T + \log^2 \delta^{-1}), & \gamma = 1, \end{cases}
\end{aligned}$$

with $N = \begin{cases} \kappa^{-1} T, & \gamma < 1 \\ \kappa^{-1} (T + \log \delta^{-1}), & \gamma = 1 \end{cases}$ steps, by taking $\gamma < \eta \lesssim 1 - T^{-1}$ when $\gamma < 1$, and $\eta = 1$ when $\gamma = 1$. In particular, in the former case, early stopping at time $T - \delta$ is not necessary, i.e. $\delta = 0$.

Proof. We have by Proposition C.4 that

$$\begin{aligned}
& \int_{s_n}^{s_{n+1}} \int_{\mathbb{X}} \left| G(\mu_s(y); \hat{\mu}_{[s_n]}^\theta(y)) - G(\mu_{s_n}(y); \hat{\mu}_{[s]}^\theta(y)) \right| \nu(dy) ds \\
& \lesssim \int_{s_n}^{s_{n+1}} \int_{\mathbb{X}} \mu_\sigma(y) \nu(dy) \\
& (\log C + \log(1 \vee (T - s_{n+1})^{-1}) + \log M) \bar{D}(s - s_n) (1 \vee (T - s_{n+1})^{-1-\gamma}) ds \\
& \lesssim (\log C + \log(1 \vee (T - s_{n+1})^{-1}) + \log M) \bar{D}^2 (s_{n+1} - s_n)^2 (1 \vee (T - s_{n+1})^{-1-\gamma}).
\end{aligned}$$

• Case 1: $\gamma < \eta \lesssim 1 - T^{-1}$

The following bound holds

$$\begin{aligned}
& \int_{s_n}^{s_{n+1}} \int_{\mathbb{X}} \left| G(\mu_s(y); \hat{\mu}_{[s_n]}^\theta(y)) - G(\mu_{s_n}(y); \hat{\mu}_{[s]}^\theta(y)) \right| \nu(dy) ds \\
& \lesssim (1 + \log(1 \vee (T - s_{n+1})^{-1})) \bar{D}^2 \kappa (s_{n+1} - s_n) (1 \vee (T - s_{n+1})^{-1-\gamma+\eta}),
\end{aligned}$$

and thus, the following error

$$\begin{aligned}
& \sum_{n=0}^{N-1} \int_{s_n}^{s_{n+1}} \int_{\mathbb{X}} \left| G(\mu_s(y); \hat{\mu}_{[s_n]}^\theta(y)) - G(\mu_{s_n}(y); \hat{\mu}_{[s]}^\theta(y)) \right| \nu(dy) ds \\
& \lesssim \sum_{n=0}^{N-1} (1 + \log(1 \vee (T - s_{n+1})^{-1})) \bar{D}^2 \kappa(s_{n+1} - s_n) (1 \vee (T - s_{n+1})^{-1-\gamma+\eta}) \\
& \lesssim \bar{D}^2 \kappa \left(T + \int_{\delta}^1 t^{-1-\gamma+\eta} \log t^{-1} dt \right) \lesssim \bar{D}^2 \kappa \left(T + \int_1^{\delta^{-1}} t^{-1-(\eta-\gamma)} \log t dt \right) \\
& \lesssim \bar{D}^2 \kappa (T + \delta^{\eta-\gamma} \log \delta^{-1}) \rightarrow \bar{D}^2 \kappa T, \quad \text{as } \delta \rightarrow 0,
\end{aligned}$$

is achievable with finite number of steps N , i.e.

$$N \lesssim \int_{\delta}^T \frac{1}{\kappa(1 \wedge t^\eta)} dt \lesssim \kappa^{-1} T + \kappa^{-1} \int_{\delta}^1 t^{-\eta} dt \lesssim \kappa^{-1} \left(T + \frac{1}{1-\eta} \right) \lesssim \kappa^{-1} T,$$

where we take $\eta \lesssim 1 - T^{-1}$.

• Case 2: $\gamma = \eta = 1$

We have the following bound

$$\begin{aligned}
& \int_{s_n}^{s_{n+1}} \int_{\mathbb{X}} \left| G(\mu_s(y); \hat{\mu}_{[s_n]}^\theta(y)) - G(\mu_{s_n}(y); \hat{\mu}_{[s]}^\theta(y)) \right| \nu(dy) ds \\
& \lesssim (1 + \log(1 \vee (T - s_{n+1})^{-1})) \bar{D}^2 \kappa(s_{n+1} - s_n) (1 \vee (T - s_{n+1})^{-1}),
\end{aligned}$$

and similarly

$$\begin{aligned}
& \sum_{n=0}^{N-1} \int_{s_n}^{s_{n+1}} \int_{\mathbb{X}} \left| G(\mu_s(y); \hat{\mu}_{[s_n]}^\theta(y)) - G(\mu_{s_n}(y); \hat{\mu}_{[s]}^\theta(y)) \right| \nu(dy) ds \\
& \lesssim \sum_{n=0}^{N-1} (1 + \log(1 \vee (T - s_{n+1})^{-1})) \bar{D}^2 \kappa(s_{n+1} - s_n) (1 \vee (T - s_{n+1})^{-1}) \\
& \lesssim \bar{D}^2 \kappa \left(T + \int_1^{\delta^{-1}} t^{-1} \log t dt \right) \lesssim \bar{D}^2 \kappa (T + \log^2 \delta^{-1}).
\end{aligned}$$

However, the number of steps N is now bounded by

$$N \lesssim \int_{\delta}^T \frac{1}{\kappa(1 \wedge t)} dt \lesssim \kappa^{-1} (T + \log \delta^{-1}).$$

□

C.4 OVERALL ERROR BOUND

Theorem C.6. *Let $\tilde{p}_{0:T}$ and $\hat{q}_{0:T}$ be the path measures of the backward process with the stochastic integral formulation (3.2) and the interpolating process (A.15) of τ -leaping algorithm (Algorithm (1)), then it holds that*

$$D_{\text{KL}}(\tilde{p}_{0:T} \| \hat{q}_{0:T}) \leq D_{\text{KL}}(\tilde{p}_0 \| \hat{q}_0) + \mathbb{E} \left[\int_0^T \int_{\mathbb{X}} \left(\mu_s(y) \log \frac{\mu_s(y)}{\hat{\mu}_{[s]}^\theta(y)} - \mu_s(y) + \hat{\mu}_{[s]}^\theta(y) \right) \nu(dy) dt \right], \quad (\text{C.2})$$

where the expectation is taken w.r.t. paths generated by the backward process (3.2).

Proof. The result directly follows by the arguments in the proof of Corollary 3.4 in Appendix A.3 and invoking the change of measure formula (Theorem 3.3) w.r.t. the true intensity μ_s as in (3.2) and the approximated intensity $\hat{\mu}_{[s]}^\theta$ as in (4.2), Proposition 4.1. □

Now, we are ready to present the proof of Theorem 4.7.

Proof of Theorem 4.7. We first rewrite the integral in (C.3) as

$$\begin{aligned} & \int_{s_n}^{s_{n+1}} \int_{\mathbb{X}} \left(\mu_s(y) \log \frac{\mu_s(y)}{\hat{\mu}_{[s]}^\theta(y)} - \mu_s(y) + \hat{\mu}_{[s]}^\theta(y) \right) \nu(dy) ds \\ &= \int_{s_n}^{s_{n+1}} \int_{\mathbb{X}} \left(\mu_{s_n}(y) \log \frac{\mu_{s_n}(y)}{\hat{\mu}_{[s_n]}^\theta(y)} - \mu_{s_n}(y) + \hat{\mu}_{[s_n]}^\theta(y) \right. \\ & \quad \left. + G(\mu_s(y); \hat{\mu}_{[s_n]}^\theta(y)) - G(\mu_{s_n}(y); \hat{\mu}_{[s_n]}^\theta(y)) \right) \nu(dy) ds. \end{aligned}$$

Therefore, the overall error is bounded by

$$\begin{aligned} & D_{\text{KL}}(\tilde{p}_{0:T-\delta} \| \hat{q}_{0:T-\delta}) \\ & \lesssim D_{\text{KL}}(\tilde{p}_0 \| \hat{q}_0) + \mathbb{E} \left[\int_0^{T-\delta} \int_{\mathbb{X}} \left(\mu_{s_n}(y) \log \frac{\mu_{s_n}(y)}{\hat{\mu}_{[s_n]}^\theta(y)} - \mu_{s_n}(y) + \hat{\mu}_{[s_n]}^\theta(y) \right) \nu(dy) dt \right] \\ & \quad + \int_{s_n}^{s_{n+1}} \int_{\mathbb{X}} \left| G(\mu_s(y); \hat{\mu}_{[s_n]}^\theta(y)) - G(\mu_{s_n}(y); \hat{\mu}_{[s_n]}^\theta(y)) \right| \nu(dy) ds \\ & \lesssim D_{\text{KL}}(p_{T-\delta} \| p_\infty) \\ & \quad + \mathbb{E} \left[\sum_{n=0}^{N-1} (s_{n+1} - s_n) \right. \\ & \quad \left. \int_{\mathbb{X}} \left(\tilde{s}_{s_n}(\tilde{x}_{s_n^-}, y) \log \frac{\tilde{s}_{s_n}(\tilde{x}_{s_n^-}, y)}{\tilde{\hat{s}}_{s_n}^\theta(\tilde{x}_{s_n^-}, y)} - \tilde{s}_{s_n}(\tilde{x}_{s_n^-}, y) + \tilde{\hat{s}}_{s_n}^\theta(\tilde{x}_{s_n^-}, y) \right) \tilde{Q}(\tilde{x}_{s_n^-}, y) \nu(dy) \right] \\ & \quad + \sum_{n=0}^{N-1} (\log C + \log M) \bar{D}^2 \kappa (s_{n+1} - s_n) \\ & \lesssim \begin{cases} \exp(-\rho T) \log |\mathbb{X}| + \epsilon + \bar{D}^2 \kappa T, & \gamma < 1, \\ \exp(-\rho T) \log |\mathbb{X}| + \epsilon + \bar{D}^2 \kappa (T + \log^2 \delta^{-1}), & \gamma = 1, \end{cases} \end{aligned}$$

where in the last inequality we used results for the first term (Truncation error, cf. Theorem C.1), the second term (Approximation error, cf. Assumption 4.6) and the third term (Discretization error, cf. Proposition C.5).

By taking

$$T = \mathcal{O} \left(\frac{\log(\epsilon^{-1} \log |\mathbb{X}|)}{\rho} \right), \quad \kappa = \mathcal{O} \left(\frac{\epsilon \rho}{\bar{D}^2 \log(\epsilon^{-1} \log |\mathbb{X}|)} \right),$$

deploying the time discretization scheme with $\gamma < \eta \lesssim 1 - T^{-1}$ when $\gamma < 1$, and $\eta = 1$ when $\gamma = 1$, and performing early stopping as

$$\delta = \begin{cases} 0, & \gamma < 1, \\ \Omega \left(\exp(-\sqrt{T}) \right), & \gamma = 1, \end{cases}$$

we have $D_{\text{KL}}(\tilde{p}_{T-\delta} \| \hat{q}_T) \leq D_{\text{KL}}(\tilde{p}_{0:T-\delta} \| \hat{q}_{0:T-\delta}) \lesssim \epsilon$ with

$$N \lesssim \kappa^{-1} T = \mathcal{O} \left(\frac{\bar{D}^2 \log^2(\epsilon^{-1} \log |\mathbb{X}|)}{\epsilon \rho^2} \right)$$

steps with probability

$$1 - \mathcal{O} \left(\kappa \bar{D}^2 T \right) = 1 - \mathcal{O}(\epsilon).$$

The proof is complete. \square

Theorem C.7. Let $\tilde{p}_{0:T}$ and $\hat{q}_{0:T}$ be the path measures of the backward process with the stochastic integral formulation (3.2) and the approximated backward process (4.4) of the uniformization algorithm (Algorithm (2)), then it holds that

$$D_{\text{KL}}(\tilde{p}_{0:T} \| q_{0:T}) \leq D_{\text{KL}}(\tilde{p}_0 \| q_0) + \mathbb{E} \left[\int_0^T \int_{\mathbb{X}} \left(\mu_s(y) \log \frac{\mu_s(y)}{\hat{\mu}_s^\theta(y)} - \mu_s(y) + \hat{\mu}_s^\theta(y) \right) \nu(dy) dt \right], \quad (\text{C.3})$$

where the expectation is taken w.r.t. paths generated by the backward process (3.2).

Proof. The result directly follows by the arguments in the proof of Corollary 3.4 in Appendix A.3 and invoking the change of measure formula (Theorem 3.3) w.r.t. the true intensity μ_s as in (3.2) and the approximated intensity $\hat{\mu}_s^\theta$ as in (4.4), Proposition 4.2. \square

Proof of Theorem 4.9. Due to the equivalence of the stochastic integral formulation (4.4) of the uniformization scheme and the approximate backward process (A.12) established in Proposition 4.2, the error for the uniformization scheme is directly bounded by the error (C.3) in Corollary 3.4, i.e.

$$\begin{aligned} & D_{\text{KL}}(\tilde{p}_{0:T-\delta} \| q_{0:T-\delta}) \\ & \leq D_{\text{KL}}(\tilde{p}_0 \| q_0) + \mathbb{E} \left[\int_0^{T-\delta} \int_{\mathbb{X}} \left(\mu_s(y) \log \frac{\mu_s(y)}{\hat{\mu}_s^\theta(y)} - \mu_s(y) + \hat{\mu}_s^\theta(y) \right) \nu(dy) dt \right] \\ & \lesssim D_{\text{KL}}(p_T \| p_\infty) \\ & + \mathbb{E} \left[\int_0^{T-\delta} \int_{\mathbb{X}} \left(\tilde{s}_s(\tilde{x}_{s-}, y) \log \frac{\tilde{s}_s(\tilde{x}_{s-}, y)}{\tilde{s}_s^\theta(\tilde{x}_{s-}, y)} - \tilde{s}_s(\tilde{x}_{s-}, y) + \tilde{s}_s^\theta(\tilde{x}_{s-}, y) \right) \tilde{Q}(\tilde{x}_{s-}, y) \nu(dy) ds \right] \\ & \lesssim |\mathbb{X}| \exp(-\rho T) + \epsilon. \end{aligned}$$

The expectation of the number of steps N is bounded by

$$\begin{aligned} \mathbb{E}[N] &= \mathbb{E} \left[\sum_{b=0}^{B-1} \mathcal{P}(\bar{\lambda}_{s_{b+1}}(s_{b+1} - s_b)) \right] = \sum_{b=0}^{B-1} \bar{\lambda}_{s_{b+1}}(s_{b+1} - s_b) \\ &\lesssim \sum_{b=0}^{B-1} \bar{D}(1 \vee (T - s_{b+1}))^{-1}(s_{b+1} - s_b) \\ &\lesssim \bar{D} \left(T + \int_\delta^1 t^{-1} dt \right) = \bar{D}(T + \log \delta^{-1}). \end{aligned}$$

By taking

$$T = \mathcal{O} \left(\frac{\log(\epsilon^{-1} \log |\mathbb{X}|)}{\rho} \right), \quad \delta = \Omega(\exp(-T))$$

we have $D_{\text{KL}}(\tilde{p}_{T-\delta} \| q_{T-\delta}) \leq D_{\text{KL}}(\tilde{p}_{0:T-\delta} \| q_{0:T-\delta}) \lesssim \epsilon$ with

$$\mathbb{E}[N] = \mathcal{O} \left(\frac{\bar{D} \log(\epsilon^{-1} \log |\mathbb{X}|)}{\rho} \right)$$

steps. \square