Convergence Guarantees of Discrete Flow Matching

Le-Tuyet-Nhi Pham
École Polytechnique, CMAP, IP Paris
Route de Saclay, 91120 Palaiseau, France
le-tuyet-nhi.pham@polytechnique.edu

Giovanni Conforti
Università degli Studi di Padova,
Via Trieste, 63, 35131 Padova, Italia
gconfort@math.unipd.it

Zhenjie Ren
LaMME, Université Évry Paris-Saclay
Boulevard François Mitterrand 91000, France
zhenjie.ren@univ-evry.fr

Alain Durmus
École Polytechnique, CMAP, IP Paris
Route de Saclay, 91120 Palaiseau, France
alain.durmus@polytechnique.edu

Abstract

Flow Matching has recently emerged as a popular class of generative models for simulating a target distribution μ_1 from samples drawn from a source distribution μ_0. This framework relies on a fixed coupling between μ_0 and μ_1, and on a deterministic or stochastic bridge to define an interpolating process between the two distributions. The time marginals of this process can then be approximately sampled by estimating the transition rates, or more generally the generator, of its Markovian projection. This framework has recently been extended to the case of discrete source and target distributions, under the name *Discrete Flow Matching* (DFM). However, theoretical guarantees for such models remain scarce. In this paper, we study two DFM models on $\mathbb{Z}_m^d = \{0, \dots, m-1\}^d$, sampled through time discretization, and derive non-asymptotic associated bounds for both of them. In contrast to previous work, we establish non-asymptotic bounds in Kullback–Leibler divergence for the early-stopped version of the target distribution. We also derive explicit convergence guarantees in total variation distance with respect to the true target distribution. Importantly, these bounds rely only on an approximation error assumption, relaxing standard score assumptions used in earlier works, while also yielding improved dependence on the vocabulary size m.

1 Introduction

Generative modeling via stochastic processes has become a central paradigm in modern machine learning as illustrated by the popularity and interest in diffusion models [SDWMG15, SE19, HJA20, SSDK+21]. Recently, Flow Matching [LCBH+23, LGL22, AVE22, ABVE25] has emerged as an alternative approach that defines a probability path between two arbitrary distributions and learns the characterization of a Markov dynamic that follows this path, resulting in simple training objectives and the ability to handle a wide variety of tasks.

More precisely, Flow Matching methods consist in simulating a target distribution μ_1 from samples drawn from a source distribution μ_0 over a finite time horizon. By employing a fixed coupling π and a reference Markov process $(R_t)_{t\in[0,1]}$, we can define an interpolant $(X_t^{\mathrm{I}})_{t\in[0,1]}$ such that $(X_0^{\mathrm{I}}, X_1^{\mathrm{I}}) \sim \pi$ and $(X_t^{\mathrm{I}})_{t\in[0,1]}$ given $(X_0^{\mathrm{I}}, X_1^{\mathrm{I}})$ has the same conditional distribution as $(R_t)_{t\in[0,1]}$ given (R_0, R_1). However, $(X_t^{\mathrm{I}})_{t\in[0,1]}$ is not Markov in general. One option then is to consider its Markov projection. Indeed, under appropriate conditions on the reference process, there exists a Markov process whose marginals are the same as the interpolated process. This process is called the Markov projection of $(X_t^{\mathrm{I}})_{t\in[0,1]}$. In particular, in the continuous setting, $(R_t)_{t\in[0,1]}$ is typically a

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

diffusion process with constant volatility, possibly zero; hence encompassing deterministic flows. Then, the resulting Markov projection has the same volatility and has an unknown drift. However, this drift is a solution of a simple regression problem and can be approximated efficiently with additional conditions on $(R_t)_{t\in[0,1]}$ [Pel23, AVE22, ABVE25, LCBH+23, Liu22, LGL22]. The same type of construction can be adapted to the discrete data. In this setting, $(R_t)_{t\in[0,1]}$ is supposed to be a continuous-time Markov chain (CTMC) and a Markov projection of $(X^{\mathrm{I}}_t)_{t\in[0,1]}$ also exists. This Markov process is also a CTMC whose generator or equivalently (concrete) score is a solution of a variational problem which can be efficiently solved under appropriate conditions on $(R_t)_{t\in[0,1]}$ [LME24, HHY+24, GRS+24, CYB+24, SGH+24, WLL+25]. This approach has shown strong empirical performance, mirroring the impact of discrete diffusion [AJH+21, HNJ+21] and score-based methods [CBB+22, SYD+23], which are now supported by many theoretical guarantees [PSO+25, BS25, CDPR25, LTSL26a].

Due to their increasing popularity, Flow Matching methods have motivated significant recent interest in understanding their theoretical foundations. While there has been a series of work providing such guarantees, e.g., [SCD24, LCHH25], in the continuous setting, achieving similar results as for continuous diffusion models, such analyses are scarcer in the discrete setting. In particular, existing results often rely on either restrictive assumptions or uniformization scheme for simulating the associated CTMC [SLH+25, WOY+25], or faces scalability issues as the vocabulary size m increases [WOX+26].

Our contribution. This work addresses these existing limitations by investigating Discrete Flow Matching (DFM) associated with two reference Markov processes: nearest-neighbor and uniform random walks. Leveraging the associated interpolants and supposing an approximation error on the estimated score, commonly assumed in the literature, we derive explicit convergence bounds of the resulting generative models adopting a Discrete Markov Probabilistic Model (DMPM) discretization scheme [PSO+25]. More specifically,

- We demonstrate that under *only* an approximation error assumption, our proposed discretization with an early-stopping rule achieves an $O(\varepsilon^2)$ error in KL divergence with respect to the early-stopped version of the target distribution with a total step complexity of $\tilde{O}(d^2/\varepsilon^2)$ for uniform random walk and $\tilde{O}(d^2m^2/\varepsilon^2)$ for nearest-neighbor random walk, where the notation $\tilde{O}$ hides logarithmic factors of d, m, ε.
- In addition, we also derive explicit convergence guarantees in total variation (TV) distance with respect to the true target distribution, which is missing in prior studies.

Our results not only improve the error's dependence on the vocabulary size m but also eliminate the restrictive boundedness conditions on the approximate score found in prior work [SLH+25, WOY+25, WOX+26]. We provide a detailed comparison of these theoretical distinctions in Section 4.

Notation. Given a measurable space $(\mathsf{E}, \mathcal{E})$, we denote by $\mathcal{P}(\mathsf{E})$ the set of probability measures of E and by 2^{E} the power set of E. Let μ, $\nu \in \mathcal{P}(\mathsf{E})$, a probability measure $\pi \in \mathcal{P}(\mathsf{E}\times\mathsf{E})$ is called a coupling of (μ,ν) if its marginals satisfy $\pi(\mathsf{A}\times\mathsf{E}) = \mu(\mathsf{A})$ and $\pi(\mathsf{E}\times\mathsf{B}) = \nu(\mathsf{B})$ for any A, $\mathsf{B} \in \mathcal{E}$. We denote by $\Pi(\mu,\nu)$ the set of all couplings of (μ,ν). We define the KL divergence (also called relative entropy) of μ with respect to ν as $\mathrm{KL}(\mu|\nu) := \int \log(\mathrm{d}\mu/\mathrm{d}\nu)\mathrm{d}\mu$ if μ is absolutely continuous with respect to ν, and $\mathrm{KL}(\mu|\nu) = +\infty$ otherwise. Given two real numbers $a, b \in \mathbb{R}$, we write $a \lesssim b$ (resp. $a \gtrsim b$) to mean $a \le Cb$ (resp. $a \ge Cb$) for a universal constant $C > 0$. Given $n \in \mathbb{N}^*$, we denote $[n] = \{1, \dots, n\}$. For any finite set A, let $|\mathsf{A}|$ denote its cardinality.

2 Discrete Flow Matching

2.1 Discrete Stochastic Interpolant and their Markov projection

We first formalize Discrete Flow Matching (DFM) by constructing Markov jump dynamics that interpolate between the two distributions, denoted by μ_0 and μ_1, over the state space $\mathbb{Z}_m^d = \{0, 1, \dots, m-1\}^d$. Our construction is similar to the original one introduced in [GRS+24]. Our presentation nevertheless differs a little taking the stochastic interpolant perspective [ABVE25] instead of the flow matching one. To this end, we first define an interpolant $(X^{\mathrm{I}}_t)_{t\in[0,1]}$ as follows.

We refer to Section A for the formal definition of the notions we used in this section and a proper introduction to CTMC.

Jump Markov Bridge and Discrete Stochastic Interpolant. Consider first a rate matrix (or generator) $(t, x, y) \in [0,1] \times (\mathbb{Z}_m^d)^2 \mapsto q_t(x, y)$. We assume $(q_t)_{t\in[0,1]}$ is a stable, conservative, and non-explosive rate matrix [FMS14]. This rate matrix defines the transition probabilities between states of the CTMC $(X_t)_{t\in[0,1]}$ with transition kernel denoted for $s, t, x_t, x_s \in ([0,1])^2 \times (\mathbb{Z}_m^d)^2$ with $s < t$ by $p_{t|s}(x_t|x_s) = \mathbb{P}(X_t = x_t|X_s = x_s)$ if $\mathbb{P}(X_s = x_s) \neq 0$ and $p_{t|s}(x_t|x_s) = \mathbb{P}(X_t = x_t)$ otherwise. Indeed, if $(q_t)_{t\in[0,1]}$ is homogeneous, i.e., $q_t = q$ for any $t \in [0,1]$, the transition probabilities of $(X_t)_{t\in[0,1]}$ are given for $s, t, x_t, x_s \in ([0,1])^2 \times (\mathbb{Z}_m^d)^2$ with $s < t$ by $p_{t|s}(x_t|x_s) = \exp((t-s)q(x_s, x_t))$. This is an easy consequence of the Forward or Backward Kolmogorov equations given in (17) and (18) respectively in the appendix. Based on this CTMC, we can consider the jump Markov bridge measure $\mathbb{B}^q$ corresponding to the conditional distribution of $(X_t)_{t\in[0,1]}$ given it starting and end points (X_0, X_1) on the space of càdlàg paths $\mathbf{D}([0,1], \mathbb{Z}_m^d)$, i.e., $\mathbb{B}^q(\cdot|X_0, X_1)$ represents the distribution of $(X_t)_{t\in[0,1]}$ given (X_0, X_1); see Section A for a formal definition. Given a coupling $\pi \in \Pi(\mu_0, \mu_1)$, we can now define an interpolated process $(X_t^{\mathrm{I}})_{t\in[0,1]}$ that bridges μ_0 and μ_1 as:

$$(X_0^{\mathrm{I}}, X_1^{\mathrm{I}}) \sim \pi\,, \quad (X_t^{\mathrm{I}})_{t\in[0,1]}\big|(X_0^{\mathrm{I}}, X_1^{\mathrm{I}}) \sim \mathbb{B}^q(\cdot|X_0^{\mathrm{I}}, X_1^{\mathrm{I}})\,. \tag{1}$$

Denote by $(p_t^{\mathrm{I}})_{t\in[0,1]}$ the time marginal densities of $(X_t^{\mathrm{I}})_{t\in[0,1]}$, that is $\mathbb{P}(X_t^{\mathrm{I}} \in \mathsf{A}) = \sum_{x\in\mathsf{A}} p_t^{\mathrm{I}}(x)$ for $t \in [0,1]$ and $\mathsf{A} \subset \mathbb{Z}_m^d$. As a straightforward consequence of the definition of $(X_t^{\mathrm{I}})_{t\in[0,1]}$, we have

$$p_t^{\mathrm{I}}(x) = \sum_{x_0,x_1\in\mathbb{Z}_m^d} p_{t|0}(x|x_0)p_{1|t}(x_1|x)\tilde{\pi}(x_0, x_1)\,, \tag{2}$$

where $\tilde{\pi}(x_0, x_1) := \pi(x_0, x_1)/p_{1|0}(x_1|x_0)$.

We consider here two reference dynamics for our framework: the Nearest-Neighbor Random Walk (NNRW) and the Uniform Random Walk (URW).

Nearest-Neighbor Random Walk. The generator associated with the NNRW is defined as:

$$q(x, y) = \begin{cases} 1/2 & \text{if } y = \sigma(x) \text{ with } \sigma \in \mathcal{M}^{\mathrm{N}}\,, \\ -d & \text{if } y = x\,, \\ 0 & \text{otherwise}\,, \end{cases} \tag{3}$$

where $\mathcal{M}^{\mathrm{N}} = \{\sigma_+^\ell, \sigma_-^\ell\}_{\ell=1}^d$ denotes the set of operators performing forward and backward jumps along the ℓ-th coordinate of the torus $\mathbb{Z}_m^d$:

$$\sigma_\pm^\ell(x) = x \pm e_\ell \pmod m\,.$$

Note that the base CTMC $(X_t)_{t\in[0,1]}$ has the component $X_t^i|X_0^i \overset{\text{i.i.d.}}{\sim} p_{t|0}^{(1)}$ for all $i \in [d]$ and $t \in [0,1]$. More precisely, for any component $i \in [d]$ and $0 \leq s < s + t \leq 1$, the displacement is given by $X_{t+s}^i - X_s^i = N_t^{i,+} - N_t^{i,-} \pmod m$, where $N_t^{i,\pm} \sim \mathrm{Poisson}(t/2)$ are independent counters for forward and backward jumps and the net displacement $S_t^i = N_t^{i,+} - N_t^{i,-}$ follows a Skellam distribution $\mathrm{Skellam}(t/2, t/2)$ [Ske46]. Thus the transition density of the full process is the product of marginals on the torus $\mathbb{Z}_m^d$:

$$p_{t+s|s}(y|x) = \prod_{i=1}^d \sum_{k\in\mathbb{Z}} \mathbb{P}(S_t^i = y^i - x^i + km)\,, \tag{4}$$

where the Skellam probability is expressed via the modified Bessel function of the first kind $I_{|a|}$:

$$\mathbb{P}(S_t^i = a) = \mathrm{e}^{-t} I_{|a|}(t)\,, \quad I_b(z) = \sum_{n=0}^\infty \frac{1}{n!(n+b)!}\left(\frac{z}{2}\right)^{2n+b} \quad \text{for } a \in \mathbb{Z},\, b \in \mathbb{N}\,.$$

The formal proof establishing the validity of (4) is detailed in Section B.1.

Uniform Random Walk. The generator associated with the URW is defined as

$$q(x,y) = \begin{cases} 1/m & \text{if } d_{\mathrm{H}}(x,y) = 1\,, \\ -d(m-1)/m & \text{if } y = x\,, \\ 0 & \text{otherwise}\,, \end{cases} \tag{5}$$

where $d_{\mathrm{H}}(x,y)$ denotes the Hamming distance between x and y. With a slight abuse of notation, we denote the set of possible jump operators as $\mathcal{M}^{\mathrm{U}} := \{\sigma_n^\ell \,:\, \ell \in [d],\, n \in [m-1]\}$, where

$$\sigma_n^\ell(x) = x + ne_\ell \pmod m$$

describes a jump of size n on the ℓ-th component. Using the properties of the Hamming distance on product spaces, the transition density $p_{s+t|s}(y|x)$ can be derived [ZCG24, Proposition 1]:

$$p_{s+t|s}(y|x) = \left(\frac{1+(m-1)\mathrm{e}^{-t}}{m}\right)^d \alpha_t^{d_{\mathrm{H}}(y,x)}\,, \quad \text{where} \quad \alpha_t := \frac{1-\mathrm{e}^{-t}}{1+(m-1)\mathrm{e}^{-t}}\,. \tag{6}$$

Markovian projection of the interpolant. A fundamental challenge is that the interpolant $(X_t^{\mathrm{I}})_{t\in[0,1]}$ is generally non-Markovian. To address this issue, we introduce its Markovian projection $(X_t^{\mathrm{M}})_{t\in[0,1]}$. As we see in our next result, this process is a Markov process uniquely characterized by a time-inhomogeneous generator that preserves the marginal distributions of the original path, that is $\mathbb{P}(X_t^{\mathrm{M}} = x) = \mathbb{P}(X_t^{\mathrm{I}} = x) = p_t^{\mathrm{I}}(x) =: \mu_t(x)$ for all $t \in [0,1]$, $x \in \mathbb{Z}_m^d$, and therefore that transports the marginals μ_0 to μ_1.

Theorem 1 (Markovian Projection of the Mixture Bridge)**.** *Let $\pi \in \Pi(\mu_0,\mu_1)$ be an arbitrary coupling and q be a homogeneous stable, conservative, and non-explosive generator. The Markovian projection $(X_t^{\mathrm{M}})_{t\in[0,1]}$ of $(X_t^{\mathrm{I}})_{t\in[0,1]}$ defined in (1) is the unique (in law) Markov process characterized by the time-inhomogeneous generator q_t^{M} such that, for all $t \in [0,1]$ and $x \neq y \in \mathbb{Z}_m^d$:*

$$q_t^{\mathrm{M}}(x,y) = q(x,y)u_t^{\mathrm{M}}(x,y)\,, \tag{7}$$

where the score function u^{M} is defined as

$$u_t^{\mathrm{M}}(x,y) = \frac{\sum_{x_0,x_1\in\mathbb{Z}_m^d} p_{1|t}(x_1|y)p_{t|0}(x|x_0)\tilde{\pi}(x_0,x_1)}{p_t^{\mathrm{I}}(x)}\,. \tag{8}$$

For completeness, the full proof of Theorem 1 is provided in Section B.2. Note that u^{M} can be expressed as the conditional expectation

$$u_t^{\mathrm{M}}(x,y) = \mathbb{E}\left[u_t^{(X_1^{\mathrm{I}})}(x,y)\middle| X_t^{\mathrm{I}} = x\right] \quad \text{where } u_t^{(x_1)}(x,y) := \frac{p_{1|t}(x_1|y)}{p_{1|t}(x_1|x)}\,. \tag{9}$$

This implies

$$q_t^{\mathrm{M}}(x,y) := \mathbb{E}\left[q_t^{(X_1^{\mathrm{I}})}(x,y) \mid X_t^{\mathrm{I}} = x\right] \quad \text{where} \quad q_t^{(x_1)}(x,y) = q(x,y)u_t^{(x_1)}(x,y)\,. \tag{10}$$

2.2 Generative Process

While the Markovian projection defines an ideal generative process, its exact simulation remains infeasible due to two main bottlenecks: (i) the intractability of the exact score defined in (8), and (ii) the requirement for time discretization in inhomogeneous CTMCs. Although techniques such as the uniformization algorithm [WOY+25] can bypass discretization, they can suffer from high computational cost and poor scalability in high-dimensional regimes. We address the former by learning an approximate score via the loss function detailed below, and the latter by employing DMPM discretization scheme [PSO+25], implementing a piecewise-constant generator. In the following sections, we formalize our generative framework, detailing the score training procedure and the construction of the resulting approximate dynamics.

Practical Training and Loss Minimization. To approximate the true discrete score u^{M}, we introduce a parameterized function class $\{u^{\mathrm{M},\theta} : [0,1] \times (\mathbb{Z}_m^d)^2 \to \mathbb{R}_+ : \theta \in \Theta\}$. Given a time-discretization $\{t_k\}_{k=0}^K$ of $[0, 1-\eta]$ associated with the step-size $h_{k+1} = t_{k+1} - t_k$ where $\eta \in (0,1)$ denotes an early-stopping parameter, the optimal parameter $\theta^\star$ is obtained by minimizing the following divergence-based objective:

$$\theta \mapsto \mathcal{L}_{\mathrm{e}}(\theta) := \sum_{k=0}^{K-1} h_{k+1} \mathbb{E}\left[\sum_{\sigma \in \mathcal{M}} q_{t_k}^{\mathrm{M},\theta} \varphi\left(\frac{q_{t_k}^{\mathrm{M}}}{q_{t_k}^{\mathrm{M},\theta}} \right) (X_{t_k}^{\mathrm{I}}, \sigma(X_{t_k}^{\mathrm{I}})) \right] , \tag{11}$$

where q^{M} is defined in (7), $q^{\mathrm{M},\theta} = q u^{\mathrm{M},\theta}$, $\varphi(a) = a\log(a) + 1 - a$ for $a \geq 0$, and $\mathcal{M}$ serves as a unified shorthand for both $\mathcal{M}^{\mathrm{N}}$ and $\mathcal{M}^{\mathrm{U}}$ for brevity. This is discrete version of the score entropy loss used in [LME24, GRS+24]. It has been already considered in the literature on theoretical guarantee for discrete diffusion models [PSO+25, LLLS25, LHL+25, DHW26]. We also consider a discretized version of the concrete score matching loss considered in [MCSE22]

$$\mathcal{L}_2(\theta) := \sum_{k=0}^{K-1} h_{k+1} \mathbb{E}\left[\sum_{\sigma \in \mathcal{M}} \left(\left(q_{t_k}^{\mathrm{M},\theta} - q_{t_k}^{\mathrm{M}} \right) (X_{t_k}^{\mathrm{I}}, \sigma(X_{t_k}^{\mathrm{I}})) \right)^2 \right] .$$

The final objective on which we will make our main approximation assumption is obtained by combining the two previous losses:

$$\mathcal{L} = \mathcal{L}_{\mathrm{e}} + \mathcal{L}_2 . \tag{12}$$

Following [MCSE22, LME24], $\mathcal{L}$ admits the following numerically tractable simplification, thanks to the conditional expectation expression of the generator in (10):

$$\tilde{\mathcal{L}}(\theta) = \sum_{k=0}^{K-1} h_{k+1} \mathbb{E}\left[\sum_{\sigma \in \mathcal{M}} \left(-q_{t_k}^{(X_1^{\mathrm{I}})} \log q_{t_k}^{\mathrm{M},\theta} + q_{t_k}^{\mathrm{M},\theta} + \left(q_{t_k}^{\mathrm{M},\theta} - q_{t_k}^{(X_1^{\mathrm{I}})} \right)^2 \right) (X_{t_k}^{\mathrm{I}}, \sigma(X_{t_k}^{\mathrm{I}})) \right] ,$$

where $q_t^{(x_1)}(x, \sigma(x))$ admits the closed-form formula in (10).

Generative Dynamics. Given the learned score $\{u_{t_k}^{\mathrm{M},\theta^\star}\}_{k=0}^K$ associated with the discretization grid $\{t_k\}_{k=0}^K$, we define the generative process as follows. Starting from the prior distribution $X_0^{\theta^\star} \sim \mu_0$, we propagate the state forward using the approximate score. For each $t \in [t_k, t_{k+1})$, given $X_{t_k}^{\theta^\star}$, the dynamics are governed by the generator $q_t^{\mathrm{M},\theta^\star}(x,y) = \tilde{u}_t^{\mathrm{M},\theta^\star}(x, y | X_{t_k}^{\theta^\star}) q(x,y)$ for $x \neq y$, where $\tilde{u}^{\mathrm{M},\theta^\star}$ is given by $\tilde{u}_t^{\mathrm{M},\theta^\star}(x, y | X_{t_k}^{\theta^\star}) = u_{t_k}^{\mathrm{M},\theta^\star}(X_{t_k}^{\theta^\star}, \sigma(X_{t_k}^{\theta^\star}))$ if $y = \sigma(x)$ for $\sigma \in \mathcal{M}$ and is 0 otherwise. Equipped with this generator, we can use an exponential distributed holding time to sample $(X_t^{\theta^\star})_{t \in [t_k, t_{k+1}]}$. This approach, referred to as DMPM, is adopted from [PSO+25]. Crucially, we stop the simulation at $t_K = 1 - \eta$ to ensure the convergence results hold despite the score's divergence at the final time step.

3 Main Results

3.1 Convergence Guarantees of DFM

To quantify the performance of DFM, we first formalize the cumulative approximation error across the training intervals.

M1 (Approximation Error)**.** There exists a constant $\tilde{\varepsilon} \geq 0$ such that

$$\mathcal{L}(\theta^\star) \leq \tilde{\varepsilon}^2 .$$

Under this requirement, we establish our main result as follows.

Theorem 2. *Consider the early-stopping parameter $\eta \in (0,1)$ and the time-discretization $\{t_k\}_{k=0}^K$ of $[0, 1-\eta]$ with $t_k = 1 - (1+h)^{-k}$, where $h \in (0,1)$ is the parameter that controls the step sizes. Under **M**1, the following holds for both NNRW and URW*

$$\mathrm{KL}(\mu_{1-\eta} | \mathrm{Law}(X_{1-\eta}^{\theta^\star})) \lesssim m^2 \eta^{-2} \tilde{\varepsilon}(\sqrt{|\mathcal{M}|} + \tilde{\varepsilon}) + (\lambda(m))^2 d^2 m^2 h \log(\eta^{-1}) \log\left(m\eta^{-1}\right) ,$$

where $\lambda(m)$ represents the fixed jump rate, taking value $1/2$ in the case of NNRW and $1/m$ for URW, and $\mathcal{M}$ denotes the collection of jump operators (see Section 2.1 for detailed definition).

Theorem 2 derives a non-asymptotic error bound for DFM under the mild assumption **M**1, explicitly decomposing the KL divergence into approximation and discretization components. We provide a sketch of the proof for this main theorem in Section 3.2, while the complete technical derivation is deferred to Section B.6.

Theorem 2 indicates that to reach a precision of $\mathrm{KL}(\mu_{1-\eta}|\mathrm{Law}(X^{\theta^\star}_{1-\eta})) = O(\varepsilon^2)$, the parameter controlling step size and the approximation error can be set as:

$$h = \frac{\varepsilon^2}{d^2 m^2 (\lambda(m))^2 \log(\eta^{-1}) \log(m\eta^{-1})} \quad \text{and} \quad \tilde{\varepsilon} = \frac{\varepsilon^2 \eta^2}{m^2 \sqrt{|\mathcal{M}|}}\,.$$

This leads to a total iteration complexity of:

$$K = \frac{\log(\eta^{-1})}{\log(1+h)} = O\left(\frac{d^2 m^2 (\lambda(m))^2}{\varepsilon^2} \log^2 \frac{1}{\eta} \log \frac{m}{\eta}\right).$$

This quadratic scaling with respect to the dimension d is consistent with existing complexity findings in [LCHH25, Corollary 5.2] and [WOX+26] . In particular, we obtain an iteration complexity of $\tilde{O}(d^2/\varepsilon^2)$ for URW and $\tilde{O}(d^2 m^2/\varepsilon^2)$ for NNRW setting (where $\tilde{O}$ omits logarithmic factors). Crucially, for the URW case, our analysis improves the dependence on the vocabulary size m compared to the recent work of [WOX+26].

Corollary 1. *Consider the early-stopping parameter $\eta \in (0,1)$ and the time-discretization $\{t_k\}_{k=0}^K$ of $[0, 1-\eta]$ with $t_k = 1-(1+h)^{-k}$, where $h \in (0,1)$ is the parameter that controls the step sizes. Under **M**1, the following holds for both NNRW and URW*

$$\|\mathrm{Law}(X^{\theta^\star}_{1-\eta}) - \mu_1\|_{\mathrm{TV}} \lesssim \mathrm{C}(m) d\eta + m\eta^{-1}\sqrt{\sqrt{|\mathcal{M}|}\tilde{\varepsilon} + \tilde{\varepsilon}^2} + \lambda(m) dm \sqrt{h \log(m\eta^{-1}) \log(\eta^{-1})},$$

where $\mathrm{C}(m)$ is defined as 1 for the URW and m for the NNRW dynamics. Here $\lambda(m)$ represents the fixed jump rate, taking value $1/2$ in the case of NNRW and $1/m$ for URW, and $\mathcal{M}$ denotes the collection of jump operators

Corollary 1 establish the explicit error bound in TV distance with respect to the true target distribution, which is missing in [WOX+26]. The proof follows directly from Theorem 2 through an application of Pinsker's and triangle inequalities. For completeness, the detailed proof of Corollary 1 is provided in Section B.7.

Remark 1. The term $\mathrm{C}(m)d\eta$ quantifies the early-stopping error, a direct consequence of truncating the generative flow at $1-\eta$ to avoid the singularity of the score at $t=1$. In the URW case, the graph is effectively complete, i.e., any state can transition to any other in a single jump. This all-to-all connectivity allows the distribution to shift globally and instantaneously, leading to a bound that is independent of the state-space size ($\mathrm{C}(m) = 1$). In contrast, the NNRW is constrained to a lattice-like topology where transitions only occur between immediate neighbors. To move probability mass across the domain, the NNRW must flow through m intermediate states, which is reflected in the linear scaling $\mathrm{C}(m) = m$. Corollary 1 implies that to achieve $O(\varepsilon)$ error in TV distance, we can set

$$\eta = \frac{\varepsilon}{\mathrm{C}(m)d}\,, \quad \tilde{\varepsilon} = \frac{\varepsilon^4}{(\mathrm{C}(m))^2 d^2 m^2 \sqrt{|\mathcal{M}|}} \quad \text{and} \quad h = \frac{\varepsilon^2}{(\lambda(m))^2 d^2 m^2 \log(m\eta^{-1}) \log(\eta^{-1})}\,.$$

This implies a total iteration complexity of:

$$K = \frac{\log(\eta^{-1})}{\log(1+h)} = \tilde{O}\left(\frac{d^2 m^2 (\lambda(m))^2}{\varepsilon^2}\right).$$

In particular, we get an iteration complexity $\tilde{O}(d^2/\varepsilon^2)$ for URW and $\tilde{O}(d^2 m^2/\varepsilon^2)$ for NNRW setting.

3.2 Proposed methodology: A Sketch of Proof

Convergence Proof. In what follows, we present a sketch of proof of the main result stated in Theorem 2 in order to outline our core methodology. To evaluate KL divergence of ground-truth against the generated distribution, we first employ the data processing inequality for the relative entropy [Nut21, Lemma 1.6]. We then decompose the total divergence into propagation mismatches

by applying Girsanov's theorem for jump processes [CDPR25, Theorem G.0.1] and substitute the formula of the piecewise-constant generator as follows

$$
\begin{aligned}
&\mathrm{KL}(\mu_{1-\eta}|\mathrm{Law}(X^{\theta^\star}_{1-\eta})) \\
&\qquad \le \mathrm{KL}(\mathrm{Law}(X^{\mathrm{M}}_{[0,1-\eta]})|\mathrm{Law}(X^{\theta^\star}_{[0,1-\eta]})) \\
&\qquad = \lambda(m)\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}} \mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left(u_t^{\mathrm{M}}\log\frac{u_t^{\mathrm{M}}}{u_{t_k}^{\mathrm{M},\theta^\star}} - u_t^{\mathrm{M}} + u_{t_k}^{\mathrm{M},\theta^\star}\right)(\sigma)\right]\mathrm{d}t \\
&\qquad \le \lambda(m)\underbrace{\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}} \mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left(\varphi(u_t^{\mathrm{M}}) - \varphi(u_{t_k}^{\mathrm{M}}) + \log u_{t_k}^{\mathrm{M}}(u_{t_k}^{\mathrm{M}} - u_t^{\mathrm{M}})\right)(\sigma)\right]\mathrm{d}t}_{E_{\text{disc}}} \\
&\qquad\quad + \underbrace{\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}} \mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left(q_t^{\mathrm{M}}\log\frac{q_{t_k}^{\mathrm{M}}}{q_{t_k}^{\mathrm{M},\theta^\star}} - q_{t_k}^{\mathrm{M}} + q_{t_k}^{\mathrm{M},\theta^\star}\right)(\sigma)\right]\mathrm{d}t}_{E_{\text{approx}}},
\end{aligned}
$$

where $\varphi(a) := a\log(a) - a + 1$ and $\lambda(m)$ represents the fixed jump rate, taking value $1/2$ for NNRW and $1/m$ for URW. Here we used the shorthand $q_t^{\mathrm{M}}(\sigma) = q_t^{\mathrm{M}}(X_t^{\mathrm{M}}, \sigma(X_t^{\mathrm{M}}))$ (and analogously for $q_{t_k}^{\mathrm{M}}(\sigma), q_{t_k}^{\mathrm{M},\theta^\star}(\sigma)$). To analyze the approximation error E_{approx}, we decompose the rate q_t^{M} as $(q_t^{\mathrm{M}}/q_{t_k}^{\mathrm{M}})q_{t_k}^{\mathrm{M}}$ to isolate the divergence term $q_{t_k}^{\mathrm{M}}\log(q_{t_k}^{\mathrm{M}}/q_{t_k}^{\mathrm{M},\theta^\star})$. We then use the following boundedness of the score to upper bound the ratio $q_t^{\mathrm{M}}/q_{t_k}^{\mathrm{M}}$ by m^2/η^2:

Lemma 1 (Uniform Boundedness of the Score)**.** *For both the NNRW and URW base processes, the score is uniformly bounded on $[0, 1-\eta]$. Specifically, for any $x \in \mathbb{Z}_m^d$, $t \in [0, 1-\eta]$ and $\sigma \in \mathcal{M}$,*

$$
\frac{1-t}{m} \lesssim u_t^{\mathrm{M}}(x, \sigma(x)) \lesssim \frac{m}{1-t}\,.
$$

Additionally, for any $x \in \mathbb{Z}_m^d$ and $t \in [0, 1-\eta]$, the following bounds hold for any $x_1 \in \mathbb{Z}_m^d$:

$$
\sum_{\sigma\in\mathcal{M}} u_t^{(x_1)}(x, \sigma(x)) \lesssim \frac{dm}{1-t} \quad \textit{and} \quad \sum_{\sigma\in\mathcal{M}} u_t^{\mathrm{M}}(x, \sigma(x)) \lesssim \frac{dm}{1-t}\,. \tag{13}
$$

Consequently, bounding E_{approx} reduces to controlling the term $q_{t_k}^{\mathrm{M}}\log(q_{t_k}^{\mathrm{M}}/q_{t_k}^{\mathrm{M},\theta^\star})$ and $q_{t_k}^{\mathrm{M},\theta^\star} - q_{t_k}^{\mathrm{M}}$, which can be addressed by applying triangle, Cauchy-Schwarz and Jensen inequalities, then using

$$
\sum_{k=0}^{K-1} h_{k+1}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left\{q_{t_k}^{\mathrm{M},\theta^\star}\varphi\left(\frac{q_{t_k}^{\mathrm{M}}}{q_{t_k}^{\mathrm{M},\theta^\star}}\right) + \left(q_{t_k}^{\mathrm{M},\theta^\star} - q_{t_k}^{\mathrm{M}}\right)^2\right\}(X_{t_k}^{\mathrm{M}}, \sigma(X_{t_k}^{\mathrm{M}}))\right] \le \tilde{\varepsilon}^2\,,
$$

which is equivalent to **M** 1 since $X_t^{\mathrm{M}} \overset{\text{dist}}{=} X_t^{\mathrm{I}}$ for all $t \in [0,1]$. We then obtain $E_{\text{approx}} = O(m^2\eta^{-2}\tilde{\varepsilon}(\sqrt{|\mathcal{M}|} + \tilde{\varepsilon}))$. To analyze E_{disc}, we first derive an ordinary differential equation (ODE) for the DFM score by leveraging the Kolmogorov equations (17) (18):

Lemma 2. *For any fixed $\sigma \in \mathcal{M}$, the score satisfies the following ODE*

$$
\partial_t u_t^{\mathrm{M}}(x, \sigma(x)) = \lambda(m)\sum_{\sigma'\in\mathcal{M}} \mathcal{A}_t^{\mathrm{M}}(x, u^{\mathrm{M}}; \sigma, \sigma')\,, \tag{14}
$$

where $\lambda(m)$ represents the fixed jump rate, taking the value $1/2$ in the case of NNRW and $1/m$ for URW. Here the operator $\mathcal{A}_t^{\mathrm{M}}$ is defined by

$$
\begin{aligned}
\mathcal{A}_t^{\mathrm{M}}(x, u^{\mathrm{M}}; \sigma, \sigma') :=& u_t^{\mathrm{M}}(x, \sigma(x))u_t^{\mathrm{M}}(x, \sigma'(x)) - u_t^{\mathrm{M}}(x, \sigma'(\sigma(x))) \\
&+ \frac{p_t^{\mathrm{I}}(\sigma'(x))}{p_t^{\mathrm{I}}(x)}\left[u_t^{\mathrm{M}}(\sigma'(x), \sigma(x)) - u_t^{\mathrm{M}}(x, \sigma(x))u_t^{\mathrm{M}}(\sigma'(x), x)\right]\,.
\end{aligned}
$$

This ODE allows us to obtain explicit evolutions for the score and its entropy function by applying Itô's formula:

Lemma 3 (Score and Entropy Dynamics). *Fix $\sigma \in \mathcal{M}$, the evolution of the score $u_t^{\mathrm{M}}(\sigma)$ and its associated entropy $\varphi(u_t^{\mathrm{M}}(\sigma))$ are governed by:*

$$\begin{aligned}
\mathbb{E}\left[u_t^{\mathrm{M}}(\sigma)\middle|\mathcal{F}_s\right] - u_s^{\mathrm{M}}(\sigma) &= \lambda(m)\int_s^t \mathbb{E}\left[\sum_{\sigma'\in\mathcal{M}} \mathcal{B}_r^{\mathrm{M}}(X_r^{\mathrm{M}}, u^{\mathrm{M}}; \sigma, \sigma')\middle|\mathcal{F}_s\right] \mathrm{d}r\,, \\
\mathbb{E}\left[\varphi(u_t^{\mathrm{M}}(\sigma))\right] - \mathbb{E}\left[\varphi(u_s^{\mathrm{M}}(\sigma))\right] &= \lambda(m)\int_s^t \mathbb{E}\left[\sum_{\sigma'\in\mathcal{M}} \mathcal{C}_r^{\mathrm{M}}(X_r^{\mathrm{M}}, u^{\mathrm{M}}; \sigma, \sigma')\right] \mathrm{d}r\,,
\end{aligned} \tag{15}$$

for any $0 \le s \le t \le 1-\eta$, where $(\mathcal{F}_t)_{t\in[0,1-\eta]}$ denotes the filtration of $(X_t^{\mathrm{M}})_{t\in[0,1-\eta]}$ and the operators $\mathcal{B}^{\mathrm{M}}, \mathcal{C}^{\mathrm{M}}$ are defined by

$$\begin{aligned}
\mathcal{B}_r^{\mathrm{M}}(x, u^{\mathrm{M}}; \sigma, \sigma') :=& u_r^{\mathrm{M}}(x, \sigma'(x)) u_r^{\mathrm{M}}(\sigma'(x), \sigma(\sigma'(x))) - u_r^{\mathrm{M}}(x, \sigma'(\sigma(x))) \\
&+ \frac{p_r^{\mathrm{I}}(\sigma'(x))}{p_r^{\mathrm{I}}(x)}\left[u_r^{\mathrm{M}}(\sigma'(x), \sigma(x)) - u_r^{\mathrm{M}}(x, \sigma(x)) u_r^{\mathrm{M}}(\sigma'(x), x)\right]\,. \\
\mathcal{C}_r^{\mathrm{M}}(x, u^{\mathrm{M}}; \sigma, \sigma') :=& u_r^{\mathrm{M}}(x, \sigma(\sigma'(x))) \log \frac{u_r^{\mathrm{M}}(\sigma'(x), \sigma(\sigma'(x)))}{u_r^{\mathrm{M}}(x, \sigma(x))} \\
&+ u_r^{\mathrm{M}}(x, \sigma'(x))\left[u_r^{\mathrm{M}}(x, \sigma(x)) - u_r^{\mathrm{M}}(\sigma'(x), \sigma(\sigma'(x)))\right]\,.
\end{aligned}$$

Substituting these expressions into discretization error gives

$$\begin{aligned}
&E_{\text{disc}} \\
&\le \lambda(m)\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\int_{t_k}^{t} \mathbb{E}\left[\sum_{\sigma,\sigma'}\left(\mathcal{C}_r^{\mathrm{M}}(X_r^{\mathrm{M}}, u^{\mathrm{M}}; \sigma, \sigma') - \log u_{t_k}^{\mathrm{M}}(\sigma)\mathcal{B}_r^{\mathrm{M}}(X_r^{\mathrm{M}}, u^{\mathrm{M}}; \sigma, \sigma')\right)\right] \mathrm{d}r\mathrm{d}t\,.
\end{aligned}$$

The primary technical challenge in the above expression lies in logarithmic terms. If the score vanishes or explodes near the boundaries of the time interval, the discretization gap becomes uncontrollable. Our analysis hinges on a key stability property shown in Lemma 1 that ensures the discrete score remains bounded on $[0, 1-\eta]$. This boundedness ensures that the logarithmic term is stable up to the early-stopping threshold η. By strategically selecting the time-grid $\{t_k\}_{k=0}^K$ to counteract the $1/(1-t)$ growth, we close the estimates and achieve the $\tilde{O}((\lambda(m))^2 d^2 m^2 h)$ total discretization error in KL divergence.

Remark 2. In standard discrete diffusion models, the score function satisfies a transitive identity: $u_t(x, z) = u_t(x, y)u_t(y, z)$, which causes the integrand in (15) to vanish, greatly simplifying the convergence analysis. In contrast, DFM lacks this property, rendering its theoretical analysis more complex than its diffusion-based counterparts.

Proof of Stability of the Score. The proof of the key Lemma 1 relies on the characterization of the projected score u_t^{M} as the conditional expectation of the bridge score $u_t^{(x_1)}$:

$$u_t^{\mathrm{M}}(x, y) = \mathbb{E}\left[u_t^{(X_1^{\mathrm{I}})}(x, y)\middle|X_t^{\mathrm{I}} = x\right]\,,$$

for $x, y \in \mathbb{Z}_m^d$ and $t \in [0, 1-\eta]$. Here, $u_t^{(x_1)}$ denotes the bridge score, which is given by the ratio of transition densities:

$$u_t^{(x_1)}(x, y) = \frac{p_{1|t}(x_1|y)}{p_{1|t}(x_1|x)}\,.$$

Consequently, the projected score inherits the uniform stability of the underlying bridge score. For the NNRW, the boundedness of $u_t^{(x_1)}$ is established by analyzing the ratios of modified Bessel functions within the Skellam density, leveraging the Amos inequality [Amo74, Eq. (9)] to bound these ratios. For the URW, stability follows directly from the explicit transition density formula, which ensures the discrete score remains bounded for all $t \in [0, 1-\eta]$. We emphasize that an *early-stopping strategy* is essential here, as the score exhibits a singularity at the terminal time $t = 1$. The detailed proof is provided in Section B.3 for completeness.

4 Related works

Discrete Diffusion Models. Discrete diffusion models (DDMs) are central to categorical generation in domains like language and chemistry [AJH+21, HNJ+21]. Recent works treat their reverse dynamics as controlled jump processes to derive non-asymptotic bounds [PSO+25, CDPR25, LLLS25, LTSL26b].

Discrete Flow Matching. Despite DFM's practical success [GRS+24, CYB+24], theoretical research remains limited. Existing analyses [SLH+25, WOY+25] typically focus on analyzing the statistical error of the one-jump-bridge-based dynamic using the exact uniformization simulation. While such algorithm eliminates discretization error theoretically, it can suffer from poor scalability. In contrast, our work implements the DMPM sampler [PSO+25] using the random-walk-based dynamics and focus on analyzing the resulting discretization error. Unlike prior studies [SLH+25, WOY+25], our framework requires no assumptions on the ground-truth or estimated generators. Furthermore, we achieve a reduction in error dependence on both the dimension d and vocabulary size m. In comparison to the recent work of [WOX+26], our analysis provides a tighter complexity bound with respect to m by leveraging the favorable properties of URW dynamic. In addition, it is worth mentioning that the assumption of the approximation error considered in [WOX+26] is different from ours. Indeed, in our notation, their condition reads:

$$\sum_{k=0}^{K-1} \mathbb{E}\left[\int_{t_k}^{t_{k+1}} \mathbb{E}\left[\sum_{z\neq X_t^{\mathrm{Alg}}} q_t^{\mathrm{M},\theta^\star} \varphi\left(\frac{q_t^{\mathrm{Alg}}}{q_t^{\mathrm{M},\theta^\star}}\right)(X_t^{\mathrm{Alg}}, z)\middle| X_{t_k}^{\mathrm{Alg}}\right] \mathrm{d}t\right] \leq \tilde{\varepsilon}^2\,, \tag{16}$$

where $(X_t^{\mathrm{Alg}})_{t\in[0,1]}$ is a process with the transition rate $(q_t^{\mathrm{Alg}})_{t\in[0,1]}$ induced by the ground-truth generator and the specific time-discretization scheme.

Their approximation assumption (16) differs substantially from **M**1. It is formulated under the law of an algorithm-generated process $(X_t^{\mathrm{Alg}})_{t\in[0,1]}$, obtained by discretizing the Markovian projection in time using a specific numerical scheme, such as τ-leaping. It controls an interval-integrated discrepancy between the exact and approximate rates along this process. Although this condition considerably simplifies the analysis, it is not directly aligned with a practical training objective. In particular, because the loss in (16) is defined under the algorithm-generated law, it is, to the best of our knowledge, intractable. Addressing this limitation is one of the main contributions of the present paper.

Our assumption **M**1, by contrast, is formulated directly in terms of a tractable, simulation-free training objective. Moreover, it admits an equivalent loss whose empirical counterpart can be evaluated using endpoint samples and samples from the explicitly known bridges of the reference process. This formulation directly connects empirical loss minimization to the approximation condition required by our convergence theorem and provides a natural basis for future analysis of generalization from the empirical loss to the population loss. Thus, the distinction between [WOX+26] and our work is not merely technical: to the best of our knowledge, ours is the *first* work to establish convergence guarantees for discrete flow-matching models under an approximation assumption directly associated with a tractable, simulation-free training objective. We summarize this comparison in Table 1.

5 Conclusion and Limitations

We establish the non-asymptotic convergence guarantees for DFM on $\mathbb{Z}_m^d$ that rigorously account for discretization error under minimal assumptions. By characterizing the score dynamics through our framework and leveraging the score's stability, we address the technical challenges posed by DFM's. We first establish convergence guarantees in KL divergence for DFM relative to the early-stopped target distribution. For URW dynamics, we achieve an $O(\varepsilon^2)$ error with an iteration complexity of $\tilde{O}(d^2/\varepsilon^2)$, significantly improving the m dependencies found in prior work. Furthermore, we derive TV distance bounds relative to the true target distribution, addressing a key gap in existing literature. However, several open questions remain. It would be valuable to complement our analysis with a statistical study of DFM for specific neural network architectures as in [SLH+25, WOY+25]. It would also be interesting to investigate whether the polynomial dependence of our convergence bounds on the early-stopping parameter can be improved.

References	Discretization scheme	Approx. assump.	Target distribution	Measure	Iteration Complexity
[WOX+26] Theorem 1	Tau-leaping and Euler	(16)	Early-stopped target $\mu_{1-\eta}$	TV distance	$\tilde{O}(d^2m/\varepsilon^2)$
[WOX+26] Theorem 3	Time-corrected sampler	(16)	Early-stopped target $\mu_{1-\eta}$	TV distance	$\tilde{O}(d^2m/\varepsilon^2)$
[WOX+26] Theorem 4	Location-corrected sampler	(16)	Early-stopped target $\mu_{1-\eta}$	TV distance	$\tilde{O}(d^2m/\varepsilon^2)$ (worst case)
This work Theorem 2	**DMPM with URW dynamic**	**M1**	**Early-stopped target $\mu_{1-\eta}$**	**KL divergence**	$\tilde{O}(d^2/\varepsilon^2)$
This work Corollary 1	**DMPM with URW dynamic**	**M1**	**True target μ_1**	**TV distance**	$\tilde{O}(d^2/\varepsilon^2)$

Table 1: Comparison of complexity results for achieving $O(\varepsilon)$ error in TV distance or $O(\varepsilon^2)$ error in KL divergence between our work and existing theoretical DFM literature. Note that all the algorithms considered are simulated only up to time $1-\eta$. The difference is that [WOX+26] establishes an error bound with respect to the *early-stopped target* $\mu_{1-\eta}$, whereas our analysis provides a convergence guarantee with respect to the *true target distribution* μ_1 (Corollary 1).

Acknowledgments and Disclosure of Funding

The work of LTNP has been supported by the École Doctorale de Mathématiques Hadamard (EDMH). AD is supported by the France 2030 program (ANR-25-PEIA-0001, THEOREM project), Hi! PARIS, and the Agence Nationale de la Recherche (Grant 11-LABX-0047). AD also received funding from the Fondation de l'École polytechnique (Servir la science campaign) and the European Union (ERC-2022-SYG-OCEAN-101071601). ZR's research is supported by the Finance For Energy Market Research Centre, the France 2030 grant (ANR-21-EXES-0003), and the PEPR PDE-AI project. Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency; neither the European Union nor the granting authority can be held responsible for them.

References

[ABVE25] Michael Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025.

[AJH+21] Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.

[Amo74] Donald E Amos. Computation of modified bessel functions and their ratios. *Mathematics of computation*, 28(125):239–251, 1974.

[AVE22] Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.

[BS25] Francis Bach and Saeed Saremi. Sampling binary data by denoising through score functions. In *Forty-second International Conference on Machine Learning*, 2025.

[CBB+22] Andrew Campbell, Joe Benton, Valentin De Bortoli, Tom Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.

[CDPR25] Giovanni Conforti, Alain Durmus, Le-Tuyet-Nhi Pham, and Gael Raoul. Non-asymptotic convergence of discrete diffusion models: Masked and random walk dynamics. *arXiv preprint arXiv:2512.00580*, 2025.

[CYB+24] 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. In *Forty-first International Conference on Machine Learning*, 2024.

[DHW26] Daniil Dmitriev, Zhihan Huang, and Yuting Wei. Efficient sampling with discrete diffusion models: Sharp and adaptive guarantees. *arXiv preprint arXiv:2602.15008*, 2026.

[Fel40] Willy Feller. On the integro-differential equations of purely discontinuous markoff processes. *Transactions of the American Mathematical Society*, 48(3):488–515, 1940.

[FMS14] Eugene A Feinberg, Manasa Mandava, and Albert N Shiryaev. On solutions of kolmogorov's equations for nonhomogeneous jump markov processes. *Journal of Mathematical Analysis and Applications*, 411(1):261–270, 2014.

[GRS+24] Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37:133345–133385, 2024.

[HHY+24] Peter Holderrieth, Marton Havasi, Jason Yim, Neta Shaul, Itai Gat, Tommi Jaakkola, Brian Karrer, Ricky TQ Chen, and Yaron Lipman. Generator matching: Generative modeling with arbitrary markov processes. *arXiv preprint arXiv:2410.20587*, 2024.

[HJA20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

[HNJ+21] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.

[Kle08] Achim Klenke. *Probability theory: a comprehensive course*. Springer, 2008.

[LCBH+23] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.

[LCHH25] Yuhao Liu, Yu Chen, Rui Hu, and Longbo Huang. Finite-time analysis of discrete-time stochastic interpolants. In *Forty-second International Conference on Machine Learning*, 2025.

[LGL22] 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.

[LHL+25] Yuchen Liang, Renxiang Huang, Lifeng Lai, Ness Shroff, and Yingbin Liang. Absorb and converge: Provable convergence guarantee for absorbing discrete diffusion models. *arXiv preprint arXiv:2506.02318*, 2025.

[Liu22] Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint arXiv:2209.14577*, 2022.

[LLLS25] Yuchen Liang, Yingbin Liang, Lifeng Lai, and Ness Shroff. Discrete diffusion models: Novel analysis and new sampler guarantees. *arXiv preprint arXiv:2509.16756*, 2025.

[LME24] Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution, 2024.

[LTSL26a] Yuchen Liang, Zhiheng Tan, Ness Shroff, and Yingbin Liang. Sharp convergence rates for masked diffusion models, 2026.

[LTSL26b] Yuchen Liang, Zhiheng Tan, Ness Shroff, and Yingbin Liang. Sharp convergence rates for masked diffusion models. *arXiv preprint arXiv:2602.22505*, 2026.

[MCSE22] 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.

[Nor98] James R Norris. *Markov chains*. Number 2. Cambridge university press, 1998.

[Nut21] Marcel Nutz. Introduction to entropic optimal transport. *Lecture notes, Columbia University*, 2021.

[Pel23] Stefano Peluchetti. Non-denoising forward-time diffusions. *arXiv preprint arXiv:2312.14589*, 2023.

[PSO+25] Le-Tuyet-Nhi Pham, Dario Shariatian, Antonio Ocello, Giovanni Conforti, and Alain Oliviero Durmus. Discrete markov probabilistic models: An improved discrete score-based framework with sharp convergence bounds under minimal assumptions. In *Forty-second International Conference on Machine Learning*, 2025.

[SCD24] Marta G Silveri, Giovanni Conforti, and Alain Durmus. Theoretical guarantees in kl for diffusion flow matching. *Advances in Neural Information Processing Systems*, 37:138432–138473, 2024.

[SDWMG15] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. pmlr, 2015.

[SE19] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.

[SGH+24] Neta Shaul, Itai Gat, Marton Havasi, Daniel Severo, Anuroop Sriram, Peter Holderrieth, Brian Karrer, Yaron Lipman, and Ricky TQ Chen. Flow matching with general discrete paths: A kinetic-optimal perspective. *arXiv preprint arXiv:2412.03487*, 2024.

[Ske46] John G Skellam. The frequency distribution of the difference between two poisson variates belonging to different populations. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 109(3):296–296, 1946.

[SLH+25] Maojiang Su, Mingcheng Lu, Jerry Yao-Chieh Hu, Shang Wu, Zhao Song, Alex Reneau, and Han Liu. A theoretical analysis of discrete flow matching generative models, 2025.

[SSDK+21] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.

[SYD+23] Haoran Sun, Lijun Yu, Bo Dai, Dale Schuurmans, and Hanjun Dai. Score-based continuous-time discrete diffusion models. In *The Eleventh International Conference on Learning Representations*, 2023.

[WLL+25] Jin Wang, Yao Lai, Aoxue Li, Shifeng Zhang, Jiacheng Sun, Ning Kang, Chengyue Wu, Zhenguo Li, and Ping Luo. Fudoki: Discrete flow-based unified understanding and generation via kinetic-optimal velocities. *arXiv preprint arXiv:2505.20147*, 2025.

[WOX+26] Zhengyan Wan, Yidong Ouyang, Liyan Xie, Fang Fang, Hongyuan Zha, and Guang Cheng. Corrected samplers for discrete flow models. *arXiv preprint arXiv:2601.22519*, 2026.

[WOY+25] Zhengyan Wan, Yidong Ouyang, Qiang Yao, Liyan Xie, Fang Fang, Hongyuan Zha, and Guang Cheng. Error analysis of discrete flow with generator matching. *arXiv preprint arXiv:2509.21906*, 2025.

[ZCG24] Zikun Zhang, Zixiang Chen, and Quanquan Gu. Convergence of score-based discrete diffusion models: A discrete-time analysis. *arXiv preprint arXiv:2410.02321*, 2024.

A Continuous Time Markov Chains and Jump Markov Bridge

Continuous Time Markov Chains. Here we present briefly Continuous-Time Markov Chains (CTMCs) as they are at the basis of the DFMs that we consider in this paper. A CTMC on $\mathbb{Z}_m^d$ is a time-indexed right-continuous stochastic process $(X_t)_{t\geq 0}$ on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ that Markov, i.e., for any $0 \leq s < t$, almost surely, $\mathbb{P}(X_{t_n} = x_{t_n}|\mathcal{F}_s) = \mathbb{P}(X_{t_n} = x_{t_n}|X_s)$, where $(\mathcal{F}_t)_{t\geq 0}$ is the natural filtration associated with $(X_t)_{t\geq 0}$.

To design CTMCs, one central object is a rate matrix, i.e., function $(x, t, \mathsf{B}) \in \mathbb{Z}_m^d \times \mathbb{R}_+ \times 2^{\mathbb{Z}_m^d} \mapsto q_t(x, \mathsf{B})$; also referred to Q-function or generator. In particular, we consider the following assumption on $(q_t)_{t\geq 0}$.

H1. $(q_t)_{t\geq 0}$ is a stable conservative rate matrix, i.e., it satisfies the following properties:

- for all $(x, t) \in \mathbb{Z}_m^d \times \mathbb{R}_+$, the function $q_t(x, \cdot)$ is a signed (discrete) measure on $\mathbb{Z}_m^d$ such that $q_t(x, \mathbb{Z}_m^d) = 0$ and $0 \leq q_t(x, \mathsf{B} \setminus \{x\}) < \infty$ for all $\mathsf{B} \subset \mathbb{Z}_m^d$;
- for all $\mathsf{B} \subset \mathbb{Z}_m^d$, the function $(x, t) \mapsto q_t(x, \mathsf{B})$ is measurable;
- for all $x \in \mathbb{Z}_m^d$, the singleton $\{x\}$ is q-bounded, i.e., $\sup_{t\in\mathbb{R}_+}(-q_t(x, \{x\})) < \infty$.

Under **H**1, the generator $(q_t)_{t\geq 0}$ allows to define a sub-Markov semigroup $\{p_{s,t} : 0 \leq s < t\}$ on $\mathbb{Z}_m^d$ and corresponding CTMCs, i.e., for $0 \leq s < t$, $p_{t|s}$ is a transition sub-probability density and the Chapman-Kolmogorov equation holds, i.e., for any $0 \leq s < u < t$ and $x_s, x_u, x_t \in \mathbb{Z}_m^d$,

$$p_{t|s}(x_t|x_s) = (p_{u|s}p_{t|u})(x_t|x_s) := \sum_{x_u\in\mathbb{Z}_m^d} p_{u|s}(x_u|x_s)p_{t|u}(x_t|x_u)\,.$$

In the case where for any $0 < s < t$, $p_{t|s} = p_{t-s|0}$, then $\{p_{s,t} : 0 \leq s < t\}$ is said to be a homogeneous semigroup. Otherwise it is said to be inhomogeneous.

Here we suppose that there is no explosion which is equivalent to the fact that the family of semigroup $\{p_{s,t} : 0 \leq s < t\}$ is in fact Markov.

H2. For any $0 \leq s < t$ and $x \in \mathbb{Z}_m^d$, $p_{s,t}(x, \mathbb{Z}_m^d) = 1$.

Remark 3. Note that [Fel40, Theorem 6] ensure that **H**2 holds if there exists a measurable function $t \mapsto \phi_t$ such that $\sup_{x\in\mathbb{Z}_m^d}(-q_t(x, \{x\})) \leq \phi_t$ for any $t \geq 0$, and for a fixed $p > 1$, $\int_s^t (\phi_u)^p \mathrm{d}u < +\infty$ for any $0 \leq s < t$.

Under **H**1 and **H**2, for any initialization $X_0 \sim \mu_0$, there exists a CTMC $(X_t)_{t\geq 0}$ starting from X_0 and associated with the family of transitions $\{p_{s,t} : 0 \leq s < t\}$:

$$p_{t|s}(x_t|x_s) = \begin{cases} \mathbb{P}(X_t = x_t|X_s = x_s) & \text{if } \mathbb{P}(X_s = x_s) \neq 0\,, \\ \mathbb{P}(X_t = x_t) & \text{otherwise}\ . \end{cases}$$

In addition, [FMS14, Theorem 4.3.] shows that for all $(x, s) \in \mathbb{Z}_m^d \times \mathbb{R}_+$ and $\mathsf{B} \subset \mathbb{Z}_m^d$ such that $\sup_{t\in\mathbb{R}_+, x\in\mathsf{B}}(-q_t(x, \{x\})) < \infty$, the function $p_{s,t}(x, \mathsf{B})$ satisfies for almost $t > s$ the forward Kolmogorov equation:

$$\frac{\partial p_{s,t}}{\partial t}(x, \mathsf{B}) = \int_{\mathsf{B}} q_t(y, \{y\})p_{s,t}(x, \mathrm{d}y) + \int_{\mathbb{Z}_m^d} q_t(y, \mathsf{B} \setminus \{y\})p_{s,t}(x, \mathrm{d}y)\,, \tag{17}$$

and the backward Kolmogorov equation:

$$\frac{\partial p_{s,t}}{\partial s}(x, \mathsf{B}) = -q_s(x, \{x\})p_{t|s}(\mathsf{B}|x) - \int_{\mathsf{X}\setminus\{x\}} q_s(x, y)p_{t|s}(\mathsf{B}|\mathrm{d}y)\,. \tag{18}$$

We note that while we restrict our presentation to a family of inhomogeneous generators $(q_t)_{t\geq 0}$ indexed by $\mathbb{R}_+$, the same results and construction apply when it is restricted to a finite interval $[0, 1]$.

Jump Markov Bridge. Based on a probability measure $\mathbb{Q}^q$ on the set of càdlàg paths $\mathbb{W} = \mathbf{D}([0,1], \mathbb{Z}_m^d)$, we define the bridge $\mathbb{B}^q$ associated with $\mathbb{Q}^q$ as the Markov kernel on $(\mathbb{Z}_m^d)^2 \times \mathbb{W}$, such that, for any measurable set $\mathsf{A} \subset \mathbb{W}$, $\mathbb{Q}^q(\mathsf{A}) = \sum_{x_0,x_1} \mathbb{B}^q(\mathsf{A}|x_0,x_1)\mathbb{Q}^q_{0,1}(x_0,x_1)$ where $\mathbb{Q}^q_{0,1}$ denotes the joint distribution of the endpoints (X_0, X_1) under $\mathbb{Q}^q$. The existence of such a kernel is guaranteed by the existence of regular conditional probabilities (see, e.g., [Kle08, Theorem 8.37]). Formally, for any finite set of time indices $\mathsf{I} = \{t_1, \dots, t_n\}$, we denote by $\mathbb{Q}^q_{\mathsf{I}}$ the I-marginal distribution of $\mathbb{Q}^q$, i.e., the pushforward measure of $\mathbb{Q}^q$ under the evaluation map $(x_t)_{t\in[0,1]} \mapsto (x_{t_1}, \dots, x_{t_n})$. From a probabilistic perspective, $\mathbb{B}^q$ represents the conditional distribution of the process $(X_t)_{t\in[0,1]}$ given its endpoints: for any bounded measurable function f on $\mathbb{W}$, $\mathbb{E}[f((X_t)_{t\in[0,1]})|X_0, X_1] = \int_{\mathbb{W}} f((x_t)_{t\in[0,1]})\mathbb{B}^q(\mathrm{d}(x_t)_{t\in[0,1]}|X_0, X_1)$.

B Technical Proofs

B.1 Proof of (4)

We first derive the expression for the one-dimensional case and subsequently extend the result to d dimensions, leveraging the mutual independence of the coordinate processes.

We denote by $q^{(1)}$ the one-dimensional generator for any coordinate process $(X^i_t)_{t\in[0,1]}$, where $i \in [d]$ and $0 \le t \le 1$, which admits the following formula for $a, b \in \mathbb{Z}_m$:

$$q^{(1)}(a,b) = \begin{cases} 1/2 & \text{if } a = b \pm 1\,, \\ -1 & \text{if } a = b\,, \\ 0 & \text{otherwise}\,. \end{cases}$$

We verify the identity $X^i_{t+s} - X^i_s = S^i_t \pmod m$ for $0 \le s < t \le 1$ by demonstrating that the kernel

$$p^{(1)}_{t+s|s}(b|a) = \sum_{k\in\mathbb{Z}} \mathbb{P}(S^i_t = b - a + km) = \mathrm{e}^{-t} \sum_{k\in\mathbb{Z}} I_{|b-a+km|}(t)$$

is the transition density for each component $i \in [d]$. This is achieved by proving that $p^{(1)}_{t+s|s}$ satisfies the one-dimensional Kolmogorov forward equation (17). Indeed, for any $a \neq b$, we have

$$\partial_t p^{(1)}_{t+s|s}(b|a) = -\mathrm{e}^{-t} \sum_{k\in\mathbb{Z}} I_{|b-a+km|}(t) + \mathrm{e}^{-t} \sum_{k\in\mathbb{Z}} \partial_t I_{|b-a+km|}(t)\,,$$

where we moved the derivative and the generator inside the summation as the sum converges absolutely and uniformly for $t > 0$. Substituting the identity $\partial_t I_n(t) = (I_{n-1}(t) + I_{n+1}(t))/2$ for $n \in \mathbb{N}^*$ yields

$$\begin{aligned} \partial_t p^{(1)}_{t+s|s}(b|a) &= -\mathrm{e}^{-t} \sum_{k\in\mathbb{Z}} I_{|b-a+km|}(t) + \frac{1}{2}\mathrm{e}^{-t} \sum_{k\in\mathbb{Z}} (I_{|b-a+km|+1} + I_{|b-a+km|-1}) \\ &= \sum_{c\in\mathbb{Z}_m} \sum_{k\in\mathbb{Z}} \mathrm{e}^{-t} I_{|c-a+km|}(t) q^{(1)}(c,b) = \sum_{c\in\mathbb{Z}_m} p^{(1)}_{t+s|s}(c|a) q^{(1)}(c,b)\,, \end{aligned}$$

where we used $I_{|z|+1} + I_{|z|-1} = I_{|z+1|} + I_{|z-1|}$ for any $z \in \mathbb{Z}^*$ in the second equality. Turning to the case where $a = b$, we obtain

$$\begin{aligned} \partial_t p^{(1)}_{t+s|s}(a|a) &= -\partial_t \left(\sum_{b\neq a} p^{(1)}_{t+s|s}(b|a) \right) = -\sum_{b\neq a} \sum_{c\in\mathbb{Z}_m} p^{(1)}_{t+s|s}(c|a) q^{(1)}(c,b) \\ &= -\sum_{c\in\mathbb{Z}_m} p^{(1)}_{t+s|s}(c|a) \sum_{b\neq a} q^{(1)}(c,b) = \sum_{c\in\mathbb{Z}_m} p^{(1)}_{t+s|s}(c|a) q^{(1)}(c,a)\,, \end{aligned}$$

where the last line follows from the fact that $\sum_{b\in\mathbb{Z}_m} q^{(1)}(c,b) = 0$ for any $c \in \mathbb{Z}_m$. The d-dimensional extension follows by showing that the product transition density $p_{t+s|s}(y|x) = \prod_{i=1}^d p^{(1)}_{t+s|s}(y^i|x^i)$ satisfies the global Kolmogorov forward equation with generator q defined in (3).

Indeed, for $x, y \in \mathbb{Z}_m^d$, we have

$$\partial_t p_{t+s|s}(y|x) = \partial_t \left(\prod_{i=1}^d p^{(1)}_{t+s|s}(y^i|x^i) \right) = \sum_{\ell=1}^d \partial_t p^{(1)}_{t+s|s}(y^\ell|x^\ell) \prod_{i \neq \ell} p^{(1)}_{t+s|s}(y^i|x^i) \,.$$

As shown before, the one-dimensional kernel $p^{(1)}_{s+t|s}$ satisfies the following for any $a, b \in \mathbb{Z}_m$:

$$\partial_t p^{(1)}_{t+s|s}(b|a) = \sum_{c \in \mathbb{Z}_m} p^{(1)}_{t+s|s}(c|a) q^{(1)}(c, b) \,,$$

which in turn implies

$$\begin{aligned}
\partial_t p_{t+s|s}(y|x) &= \sum_{\ell=1}^d \sum_{j \in \mathbb{Z}_m} p^{(1)}_{t+s|s}(j|x^\ell) q^{(1)}(j, y^\ell) \prod_{i \neq \ell} p^{(1)}_{t+s|s}(y^i|x^i) \\
&= \sum_{\ell=1}^d \sum_{j=-1}^1 p^{(1)}_{t+s|s}(y^\ell + j|x^\ell) q^{(1)}(y^\ell + j, y^\ell) \prod_{i \neq \ell} p^{(1)}_{t+s|s}(y^i|x^i) \\
&= \sum_{\sigma \in \mathcal{M}} p_{t+s|s}(\sigma(y)|x) q(\sigma(y), y) + p_{t+s|s}(y|x) q(y, y) \\
&= \sum_{z \in \mathbb{Z}_m^d} p_{t+s|s}(z|x) q(z, y) \,.
\end{aligned}$$

Consequently, $p_{t+s|s}$ satisfies the Kolmogorov forward equation associated with the nearest-neighbor random walk generator q. Invoking the uniqueness of solutions to the Kolmogorov equation [Nor98, Theorem 2.8.3], we conclude that $p_{t+s|s}$ is the true transition density of $(X_t)_{t \in [0,1]}$, thereby validating its representation via the Poisson process.

B.2 Proof of Theorem 1: Markovian Projection

For fixed $x_0, x_1 \in \mathbb{Z}_m^d$ and $0 \leq t < t + \Delta t \leq 1$, we have

$$\mathbb{P}(X_{t+\Delta t} = y \mid X_t = x,\ X_1 = x_1) = \frac{p_{t+\Delta t|t}(y|x) p_{1|t+\Delta t}(x_1|y)}{p_{1|t}(x_1|x)} \,.$$

Recall that by Kolmogorov equation, $p_{t+\Delta t|t}(y|x) = \mathbf{1}_{y=x} + q(x, y)\Delta t + o(\Delta t)$ with q is the homogeneous generator of $(X_t)_{t \in [0,1]}$ and analogously for the bridge generator $q_t^{x_0 \to x_1}$ between x_0 and x_1. Plugging this into the preceding equation yields

$$q_t^{x_0 \to x_1}(x, y)\Delta t + o(\Delta t) = q(x, y) \frac{p_{1|t+\Delta t}(x_1|y)}{p_{1|t}(x_1|x)} \Delta t + o(\Delta t) \,,$$

for $x \neq y$. Divide both hand sides by Δt and let $\Delta t \to 0$, we obtain that

$$q_t^{x_0 \to x_1}(x, y) = q(x, y) \frac{p_{1|t}(x_1|y)}{p_{1|t}(x_1|x)} \,.$$

Note that because the expression for $q_t^{x_0 \to x_1}$ is independent of the starting point x_0, it can be written more simply as $q_t^{(x_1)}$. Using this notation, we define the transition rate for $x \neq y$ over the interval $t \in [0, 1]$ as:

$$q_t^{\mathrm{M}}(x, y) \ := \ \mathbb{E}\left[q_t^{(X_1^{\mathrm{I}})}(x, y) \,\middle|\, X_t^{\mathrm{I}} = x \right] \,,$$

with the diagonal terms $q_t^{\mathrm{M}}(x, x) = -\sum_{y \neq x} q_t^{\mathrm{M}}(x, y)$ chosen to enforce the row-sum constraint. A direct calculation yields the explicit form of the generator $q_t^{\mathrm{M}}(x, y)$ for $y \neq x$:

$$q_t^{\mathrm{M}}(x, y) = q(x, y) u_t^{\mathrm{M}}(x, y) \,,$$

where the discrete score function u^{M} is defined as

$$u_t^{\mathrm{M}}(x,y) := \frac{\sum_{x_0,x_1\in\mathbb{Z}_m^d} p_{1|t}(x_1|y)p_{t|0}(x|x_0)\tilde{\pi}(x_0,x_1)}{p_t^{\mathrm{I}}(x)}\,,$$

where $\tilde{\pi}(x_0,x_1) = \pi(x_0,x_1)/p_{1|0}(x_1|x_0)$. We now verify that the projected generator q_t^{M} induces a probability flow that preserves the marginals of the mixture bridge $(X_t^{\mathrm{I}})_{t\in[0,1]}$ by showing that p_t^{I} satisfies the Kolmogorov forward equation associated with q_t^{M}. Recall that $(s,t,x,y)\mapsto p_{t|s}(y|s)$ satisfies for any $x,y\in\mathbb{Z}_m^d$ and $0\le s<t\le 1$ both the Kolmogorov forward and backward equation (17), (18). Exploiting these equations and (2), for $t\in(0,1)$, we get that

$$\begin{aligned}
\partial_t p_t^{\mathrm{I}}(x) &= \sum_{x_0,x_1\in\mathbb{Z}_m^d} \tilde{\pi}(x_0,x_1)\left[\partial_t p_{t|0}(x|x_0)p_{1|t}(x_1|x) + p_{t|0}(x|x_0)\partial_t p_{1|t}(x_1|x)\right]\\
&= \sum_{x_0,x_1\in\mathbb{Z}_m^d} \tilde{\pi}(x_0,x_1)\sum_{z\in\mathbb{Z}_m^d}\left[p_{t|0}(z|x_0)q(z,x)p_{1|t}(x_1|x) - q(x,z)p_{1|t}(x_1|z)p_{t|0}(x|x_0)\right]\\
&= \sum_{x_0,x_1\in\mathbb{Z}_m^d} \tilde{\pi}(x_0,x_1)\sum_{z\neq x}\left[p_{t|0}(z|x_0)q(z,x)p_{1|t}(x_1|x) - q(x,z)p_{1|t}(x_1|z)p_{t|0}(x|x_0)\right]\\
&= \sum_{z\neq x} p_t^{\mathrm{I}}(z)q_t^{\mathrm{M}}(z,x) - p_t^{\mathrm{I}}(x)\underbrace{\sum_{z\neq x} q_t^{\mathrm{M}}(x,z)}_{-q_t^{\mathrm{M}}(x,x)} = \sum_{z\in\mathbb{Z}_m^d} p_t^{\mathrm{I}}(z)q_t^{\mathrm{M}}(z,x)\,.
\end{aligned}$$

Thus, the process $(X_t^{\mathrm{M}})_{t\in[0,1]}$ induced by q^{M} matches the marginals of the mixture bridge $(X_t^{\mathrm{I}})_{t\in[0,1]}$. Since $(X_t^{\mathrm{M}})_{t\in[0,1]}$ is Markovian by construction, and given that the linearity of the Kolmogorov forward equation ensures the uniqueness of the generator q^{M} within the support of the marginals, the Markovian projection is unique in law. This concludes the proof of the results in Theorem 1.

B.3 Proof of Lemma 1: Stability of the Score

Recall that for any $x\neq y$ and $t\in[0,1-\eta]$,

$$u_t^{\mathrm{M}}(x,y) = \mathbb{E}\left[u_t^{(X_1^{\mathrm{I}})}(x,y)\middle| X_t^{\mathrm{I}} = x\right]\,,$$

where $u_t^{(x_1)}$ denotes the bridge score and admits the following formula

$$u_t^{(x_1)}(x,y) = \frac{p_{1|t}(x_1|y)}{p_{1|t}(x_1|x)}\,.$$

Therefore, it suffices to bound $u_t^{(x_1)}(x,y)$ to get the uniform bound of the score $u_t^{\mathrm{M}}(x,y)$. As $q(x,y)=0$ for $y\notin\{x,\sigma(x):\sigma\in\mathcal{M}\}$, it suffices to focus on the case $y=\sigma(x)$ with $\sigma\in\mathcal{M}$. We now evaluate the NNRW and URW cases individually to establish their respective bounds.

Nearest-Neighbor Random Walk. We begin by recalling

$$u_t^{(x_1)}(x,\sigma_\pm^\ell(x)) = \frac{p_{1|t}(x_1|\sigma_\pm^\ell(x))}{p_{1|t}(x_1|x)} \overset{(4)}{=} \frac{p_{1|t}^{(1)}(x_1^\ell|x^\ell\pm 1)}{p_{1|t}^{(1)}(x_1^\ell|x^\ell)} = \frac{\sum_{k\in\mathbb{Z}} I_{|\Delta\pm 1+km|}(1-t)}{\sum_{k\in\mathbb{Z}} I_{|\Delta+km|}(1-t)}\,,$$

where $\Delta := x_1^\ell - x^\ell$. Let us investigate $\sum_{k\in\mathbb{Z}} I_{|a+km|}(t)$ for $a\in\mathbb{Z}$, $t\in[0,1-\eta]$ to derive the bound on the interpolated score $u_t^{(x_1)}(x,\sigma_\pm^\ell(x))$. Let

$$\tilde{a} = \min_{k\in\mathbb{Z}} |a+km| \le \min_{k\in\mathbb{Z}}(|a|+|k|m) \le \min_{k\in\mathbb{Z}}(1+|k|)m - 1 = m-1\,.$$

We write

$$\sum_{k\in\mathbb{Z}} I_{|a+km|}(t) = I_{\tilde{a}}(t)\sum_{k\in\mathbb{Z}}\frac{I_{|a+km|}(t)}{I_{\tilde{a}}(t)}\,. \tag{19}$$

For $|a+km|-\tilde{a}\geq 1$, we have

$$\frac{I_{|a+km|}(t)}{I_{\tilde{a}}(t)}=\prod_{j=0}^{|a+km|-\tilde{a}-1}\frac{I_{\tilde{a}+j+1}(t)}{I_{\tilde{a}+j}(t)}=\frac{I_{\tilde{a}+1}(t)}{I_{\tilde{a}}(t)}\prod_{j=1}^{|a+km|-\tilde{a}-1}\frac{I_{\tilde{a}+j+1}(t)}{I_{\tilde{a}+j}(t)}\,.$$

Note that $I_b(t)$ is decreasing in $b\in\mathbb{Z}^+$ for fixed $t\in[0,1-\eta]$, thus $I_{\tilde{a}+1}(t)\leq I_{\tilde{a}}(t)$ and

$$\frac{I_{|a+km|}(t)}{I_{\tilde{a}}(t)}\leq\prod_{j=1}^{|a+km|-\tilde{a}-1}\frac{I_{\tilde{a}+j+1}(t)}{I_{\tilde{a}+j}(t)}=\prod_{j=1}^{|a+km|-\tilde{a}-1}R^+_{\tilde{a}+j}(t)\,,$$

where

$$R^+_b(t):=\frac{I_{b+1}(t)}{I_b(t)}\quad\text{for } b\in\mathbb{Z}^+\,.$$

By Amos inequality [Amo74, Equation (9)], we have

$$R^+_{\tilde{a}+j}(t)\leq\frac{t}{\tilde{a}+j+\sqrt{t^2+(\tilde{a}+j)^2}}\leq\frac{t}{2(\tilde{a}+j)}\overset{j\geq 1}{\leq}\frac{t}{2(\tilde{a}+1)}\,.$$

Hence for $|a+km|-\tilde{a}\geq 1$,

$$\frac{I_{|a+km|}(t)}{I_{\tilde{a}}(t)}\leq\prod_{j=1}^{|a+km|-\tilde{a}-1}\frac{t}{2(\tilde{a}+1)}=\left(\frac{t}{2(\tilde{a}+1)}\right)^{|a+km|-\tilde{a}-1}\,.$$

Plugging it into (19) gives

$$\begin{aligned}\sum_{k\in\mathbb{Z}}I_{|a+km|}(t)\leq&I_{\tilde{a}}(t)\left[1+\sum_{\substack{k\in\mathbb{Z}\\|a+km|-\tilde{a}\geq 1}}\left(\frac{t}{2(\tilde{a}+1)}\right)^{|a+km|-\tilde{a}-1}\right]\\\leq&I_{\tilde{a}}(t)\left[1+2\sum_{k=0}^{\infty}\left(\frac{t}{2(\tilde{a}+1)}\right)^k\right]=I_{\tilde{a}}(t)\left[1+\frac{2}{1-\frac{t}{2(\tilde{a}+1)}}\right]\\\leq&I_{\tilde{a}}(t)\Big[1+\underbrace{\frac{4(\tilde{a}+1)}{2(\tilde{a}+1)-1}}_{\leq 4}\Big]\leq 5I_{\tilde{a}}(t)\,.\end{aligned}$$

This together with the lower bound $\sum_{k\in\mathbb{Z}}I_{|a+km|}(t)\geq I_{\tilde{a}}(t)$ imply

$$I_{\tilde{a}}(t)\leq\sum_{k\in\mathbb{Z}}I_{|a+km|}(t)\lesssim I_{\tilde{a}}(t)\,.$$

As a result,

$$\underbrace{\frac{I_{\min_k|\Delta+km\pm 1|}(1-t)}{I_{\min_k|\Delta+km|}(1-t)}}_{R}\lesssim\underbrace{\frac{\sum_{k\in\mathbb{Z}}I_{|\Delta\pm 1+km|}(1-t)}{\sum_{k\in\mathbb{Z}}I_{|\Delta+km|}(1-t)}}_{=u_t^{(x_1)}(x,\sigma^\ell_\pm(x))}\lesssim\frac{I_{\min_k|\Delta+km\pm 1|}(1-t)}{I_{\min_k|\Delta+km|}(1-t)}$$

We evaluate the ratio R by leveraging the decreasing property of $I_b(1-t)$ in $b\in\mathbb{Z}^+$ for fixed $t\in[0,1]$ as follows

$$\begin{aligned}R\geq&\frac{I_{\min_k|\Delta+km|+1}(1-t)}{I_{\min_k|\Delta+km|}(1-t)}=R^+_{\min_k|\Delta+km|}(1-t)\\\geq&\frac{1-t}{\min_k|\Delta+km|+1+\sqrt{(1-t)^2+(\min_k|\Delta+km|+1)^2}}\\\geq&\frac{1-t}{m+\sqrt{(1-t)^2+m^2}}\geq\frac{1-t}{2m+1}\gtrsim\frac{1-t}{m}\,,\end{aligned}$$

where we used Amos inequality [Amo74, Equation (9)] in the second line. Analogously, we can obtain the upper bound of R:

$$R \leq \frac{I_{\min_k |\Delta+km\pm 1|}(1-t)}{I_{\min_k |\Delta+km\pm 1|+1}(1-t)} = \frac{1}{R^+_{\min_k |\Delta+km\pm 1|}(1-t)} \lesssim \frac{m}{1-t}\,.$$

Altogether yields the final bound on $u_t^{(x_1)}(x, \sigma^\ell_\pm(x))$:

$$\frac{1-t}{m} \lesssim u_t^{(x_1)}(x, \sigma^\ell_\pm(x)) \lesssim \frac{m}{1-t}\,.$$

By taking the conditional expectation, we deduce the same uniform bound on $u_t^{\mathrm{M}}(x, \sigma^\ell_\pm(x))$:

$$\frac{1-t}{m} \lesssim u_t^{\mathrm{M}}(x, \sigma^\ell_\pm(x)) \lesssim \frac{m}{1-t}\,.$$

Moreover, the bound of the bridge score also implies

$$\sum_{\sigma\in\mathcal{M}} u_t^{(x_1)}(x, \sigma(x)) = \sum_{\ell=1}^{d} u_t^{(x_1)}(x, \sigma^\ell_+(x)) + \sum_{\ell=1}^{d} u_t^{(x_1)}(x, \sigma^\ell_-(x)) \lesssim \frac{dm}{1-t}\,.$$

Taking the conditional expectation leads to an identical bound for $\sum_{\sigma\in\mathcal{M}} u_t^{\mathrm{M}}(x, \sigma(x))$.

Uniform Random Walk. We begin by recalling

$$u_t^{(x_1)}(x,y) = \frac{p_{1|t}(x_1|y)}{p_{1|t}(x_1|x)} \overset{(6)}{=} \frac{\alpha_{1-t}^{d_{\mathrm{H}}(y,x_1)}}{\alpha_{1-t}^{d_{\mathrm{H}}(x,x_1)}} = \alpha_{1-t}^{d_{\mathrm{H}}(y,x_1)-d_{\mathrm{H}}(x,x_1)}\,. \tag{20}$$

By triangle inequality, we get that

$$d_{\mathrm{H}}(x,x_1) - d_{\mathrm{H}}(y,x) \leq d_{\mathrm{H}}(y,x_1) \leq d_{\mathrm{H}}(x,x_1) + d_{\mathrm{H}}(x,y)\,,$$

which implies the following estimate since $\alpha_t < 1$ for any $t \in [0, 1-\eta]$:

$$\alpha_{1-t}^{-d_{\mathrm{H}}(y,x)} \geq u_t^{(x_1)}(x,y) \geq \alpha_{1-t}^{d_{\mathrm{H}}(y,x)}\,. \tag{21}$$

For $y = \sigma_n^\ell(x)$ with $\ell \in [d], n \in [m-1]$, this estimate reduces to

$$\alpha_{1-t}^{-1} \geq u_t^{(x_1)}(x, \sigma_n^\ell(x)) \geq \alpha_{1-t}\,,$$

which yields the same bound on $u_t^{\mathrm{M}}(x, \sigma_n^\ell(x))$. Note that α_{1-t}^{-1} admits the following upper bound

$$\alpha_{1-t}^{-1} = \frac{1+(m-1)\mathrm{e}^{-(1-t)}}{1-\mathrm{e}^{-(1-t)}} = \frac{\mathrm{e}^{1-t}+m-1}{\mathrm{e}^{1-t}-1} \leq \frac{\mathrm{e}+m-1}{1-t} \lesssim \frac{m}{1-t}\,, \tag{22}$$

where we used the inequality $\mathrm{e}^x \geq x+1$ for any $x \in \mathbb{R}$ together with the fact $m \geq 2$. Thus

$$\frac{1-t}{m} \lesssim u_t^{\mathrm{M}}(x, \sigma_n^\ell(x)) \lesssim \frac{m}{1-t}\,.$$

Furthermore, (20) implies

$$\begin{aligned}
\sum_{\sigma\in\mathcal{M}} u_t^{(x_1)}(x,\sigma(x)) &= \sum_{\ell=1}^{d}\sum_{n=1}^{m-1} \alpha_{1-t}^{d_{\mathrm{H}}(x_1,\sigma_n^\ell(x))-d_{\mathrm{H}}(x_1,x)} \\
&= \sum_{\ell=1}^{d}\sum_{n=1}^{m-1} \Big(\underbrace{\alpha_{1-t}}_{\leq 1} \mathbf{1}_{x_1^\ell = x^\ell} + \alpha_{1-t}^{-1}\mathbf{1}_{x_1^\ell = (\sigma_n^\ell(x))^\ell} + \mathbf{1}_{x_1^\ell \notin \{x^\ell, (\sigma_n^\ell(x))^\ell\}} \Big) \\
&\leq \sum_{\ell=1}^{d} \left(\alpha_{1-t}^{-1} + \sum_{n=1}^{m-1} \mathbf{1}_{x_1^\ell \neq (\sigma_n^\ell(x))^\ell} \right) \\
&\overset{(22)}{\lesssim} \sum_{\ell=1}^{d} \left(\frac{m}{1-t} + m - 1 \right) \lesssim \frac{dm}{1-t}\,.
\end{aligned}$$

Taking the conditional expectation leads to an identical bound for $\sum_{\sigma\in\mathcal{M}} u_t^{\mathrm{M}}(x, \sigma(x))$. The proof of Lemma 1 is complete.

B.4 Proof of Lemma 2

Fix $\sigma \in \mathcal{M}$ and recall the following formula of the discrete score u^{M} for $x \in \mathbb{Z}_m^d$, $t \in [0, 1-\eta]$:

$$u_t^{\mathrm{M}}(x, \sigma(x)) = \frac{\sum_{x_0, x_1 \in \mathbb{Z}_m^d} p_{t|0}(x|x_0) p_{1|t}(x_1|y) \tilde{\pi}(x_0, x_1)}{p_t^{\mathrm{I}}(x)}, .$$

By quotient rule, we get that

$$\begin{aligned}\partial_t u_t^{\mathrm{M}}(x, \sigma(x)) = &\sum_{x_0, x_1 \in \mathbb{Z}_m^d} \frac{\tilde{\pi}(x_0, x_1)}{p_t^{\mathrm{I}}(x)} \left[p_{1|t}(x_1|\sigma(x)) \partial_t p_{t|0}(x|x_0) + p_{t|0}(x|x_0) \partial_t p_{1|t}(x_1|\sigma(x)) \right] \\ &- \sum_{x_0, x_1 \in \mathbb{Z}_m^d} \frac{\tilde{\pi}(x_0, x_1)}{(p_t^{\mathrm{I}}(x))^2} p_{t|0}(x|x_0) p_{1|t}(x_1|\sigma(x)) \partial_t p_t^{\mathrm{I}}(x) .\end{aligned}$$

Employing the Kolmogorov forward and backward equations (17), (18) satisfied by $p_{t|0}$, $p_{1|t}$ and p_t^{I}, we arrive at

$$\begin{aligned}&\partial_t u_t^{\mathrm{M}}(x, \sigma(x)) \\ &= \lambda(m) \sum_{\sigma' \in \mathcal{M}} \sum_{x_0, x_1 \in \mathbb{Z}_m^d} \frac{\tilde{\pi}(x_0, x_1)}{p_t^{\mathrm{I}}(x)} \Bigg[p_{1|t}(x_1|\sigma(x)) p_{t|0}(\sigma'(x)|x_0) - p_{1|t}(x_1|\sigma(x)) p_{t|0}(x|x_0) \\ &\qquad - p_{t|0}(x|x_0) p_{1|t}(x_1|\sigma'(\sigma(x))) + p_{1|t}(x_1|\sigma(x)) p_{t|0}(x|x_0) \\ &\qquad + p_{t|0}(x|x_0) p_{1|t}(x_1|\sigma(x)) \left(u_t^{\mathrm{M}}(x, \sigma'(x)) - \frac{p_t^{\mathrm{I}}(\sigma'(x))}{p_t^{\mathrm{I}}(x)} u_t^{\mathrm{M}}(\sigma'(x), x) \right) \Bigg] \\ &= \lambda(m) \sum_{\sigma' \in \mathcal{M}} \Bigg\{ u_t^{\mathrm{M}}(x, \sigma'(x)) u_t^{\mathrm{M}}(x, \sigma(x)) - u_t^{\mathrm{M}}(x, \sigma'(\sigma(x))) \\ &\qquad + \frac{p_t^{\mathrm{I}}(\sigma'(x))}{p_t^{\mathrm{I}}(x)} \left[u_t^{\mathrm{M}}(\sigma'(x), \sigma(x)) - u_t^{\mathrm{M}}(\sigma'(x), x) u_t^{\mathrm{M}}(x, \sigma(x)) \right] \Bigg\},\end{aligned}$$

where $\lambda(m)$ represents the fixed jump rate, taking the value $1/2$ in the case of NNRW and $1/m$ for URW. Denoting the operator $\mathcal{A}^{\mathrm{M}}$ as in Lemma 2 concludes the proof.

B.5 Proof of Lemma 3: Score and Entropy Evolution

Score Evolution. Fixing $\sigma \in \mathcal{M}$, we apply Itô's formula to the process $f^\sigma(t, X_t^{\mathrm{M}})$, where $f^\sigma(t, x) := u_t^{\mathrm{M}}(x, \sigma(x))$. Observing that $X_t^{\mathrm{M}} = X_{t-}^{\mathrm{M}}$ for Lebesgue-almost every $t \in [0, 1-\eta]$, we find that

$$f^\sigma(t, X_t^{\mathrm{M}}) - f^\sigma(s, X_s^{\mathrm{M}}) - \int_s^t \left[\partial_r f^\sigma(r, X_r^{\mathrm{M}}) + (q^{\mathrm{M}} f^\sigma)(r, X_r^{\mathrm{M}}) \right] \mathrm{d}r$$

is a local martingale on $0 \le s < t \le 1-\eta$. In view of the finite state space and the fact that u_t^{M} is uniformly bounded for all $t \in [0, 1-\eta]$ (as shown in Lemma 1), we conclude that this local martingale is integrable and thus a true martingale. Denoting by $b_r^\sigma(X_r^{\mathrm{M}})$ the integrand above, we

have

$$
\begin{aligned}
b_r^\sigma(x) &= \partial_r u_r^{\mathrm{M}}(x,\sigma(x)) + \lambda(m)\sum_{\sigma'\in\mathcal{M}}\left[u_r^{\mathrm{M}}(\sigma'(x),\sigma(\sigma'(x))) - u_r^{\mathrm{M}}(x,\sigma(x))\right] u_r^{\mathrm{M}}(x,\sigma'(x)) \\
&\overset{(14)}{=} \lambda(m)\sum_{\sigma'\in\mathcal{M}}\Bigg\{u_r^{\mathrm{M}}(x,\sigma(x))u_r^{\mathrm{M}}(x,\sigma'(x)) - u_r^{\mathrm{M}}(x,\sigma'(\sigma(x))) \\
&\qquad\qquad + \frac{p_r^{\mathrm{I}}(\sigma'(x))}{p_r^{\mathrm{I}}(x)}\left[u_r^{\mathrm{M}}(\sigma'(x),\sigma(x)) - u_r^{\mathrm{M}}(x,\sigma(x))u_r^{\mathrm{M}}(\sigma'(x),x)\right] \\
&\qquad\qquad + \left[u_r^{\mathrm{M}}(\sigma'(x),\sigma(\sigma'(x))) - u_r^{\mathrm{M}}(x,\sigma(x))\right] u_r^{\mathrm{M}}(x,\sigma'(x))\Bigg\} \\
&= \lambda(m)\sum_{\sigma'\in\mathcal{M}}\Bigg\{u_r^{\mathrm{M}}(x,\sigma'(x))u_r^{\mathrm{M}}(\sigma'(x),\sigma(\sigma'(x))) - u_r^{\mathrm{M}}(x,\sigma'(\sigma(x))) \\
&\qquad\qquad + \frac{p_r^{\mathrm{I}}(\sigma'(x))}{p_r^{\mathrm{I}}(x)}\left[u_r^{\mathrm{M}}(\sigma'(x),\sigma(x)) - u_r^{\mathrm{M}}(x,\sigma(x))u_r^{\mathrm{M}}(\sigma'(x),x)\right]\Bigg\}.
\end{aligned}
$$

Consequently, we obtain the desired conclusion: for $0 \le s < t \le 1-\eta$,

$$
\mathbb{E}\left[u_t^{\mathrm{M}}(\sigma)\middle|\mathcal{F}_s\right] - u_s^{\mathrm{M}}(\sigma) = \lambda(m)\int_s^t \mathbb{E}\left[\sum_{\sigma'\in\mathcal{M}}\mathcal{B}_r^{\mathrm{M}}(X_r^{\mathrm{M}},u^{\mathrm{M}};\sigma,\sigma')\middle|\mathcal{F}_s\right]\mathrm{d}r\,,
$$

where $(\mathcal{F}_t)_{t\in[0,1-\eta]}$ denotes the filtration of $(X_t^{\mathrm{M}})_{t\in[0,1-\eta]}$ and the operator $\mathcal{B}^{\mathrm{M}}$ defined as in Lemma 3.

Entropy Score Evolution. Next, we investigate the evolution of $\sum_{\sigma\in\mathcal{M}}\varphi(u_t^{\mathrm{M}}(\sigma))$-a central component in the convergence derivation, where $\varphi(a) := a\log(a) - a + 1$. Analogously as before, by Itô's formula, the following process

$$
\varphi(u_t^{\mathrm{M}}(\sigma)) - \varphi(u_s^{\mathrm{M}}(\sigma)) - \int_s^t\left[\partial_r\varphi(u_r^{\mathrm{M}}(\sigma)) + (q_r^{\mathrm{M}}\varphi(u_r^{\mathrm{M}}))(X_r^{\mathrm{M}},\sigma(X_r^{\mathrm{M}}))\right]\mathrm{d}r
$$

is a local martingale for $0 \le s < t \le 1-\eta$. The finiteness of the state space, coupled with the uniform bound on u_t^{M} provided in Lemma 1, ensures that the local martingale is integrable. Consequently, it becomes a true martingale throughout $[0, 1-\eta]$. Denoting by $c_r^\sigma(X_r^{\mathrm{M}})$ the above integrand, we have

$$
\begin{aligned}
&c_r^\sigma(x) \\
&= \partial_r\varphi(u_r^{\mathrm{M}}(x,\sigma(x))) + \lambda(m)\sum_{\sigma'\in\mathcal{M}}\left[\varphi(u_r^{\mathrm{M}}(\sigma'(x),\sigma(\sigma'(x)))) - \varphi(u_r^{\mathrm{M}}(x,\sigma(x)))\right] u_r^{\mathrm{M}}(x,\sigma'(x)) \\
&= \log u_r^{\mathrm{M}}(x,\sigma(x))\partial_r u_r^{\mathrm{M}}(x,\sigma(x)) + \lambda(m)\sum_{\sigma'\in\mathcal{M}} u_r^{\mathrm{M}}(x,\sigma'(x)) \\
&\qquad\qquad \left[(u_r^{\mathrm{M}}\log u_r^{\mathrm{M}} - u_r^{\mathrm{M}})(\sigma'(x),\sigma(\sigma'(x))) - (u_r^{\mathrm{M}}\log u_r^{\mathrm{M}} - u_r^{\mathrm{M}})(x,\sigma(x))\right].
\end{aligned}
$$

Substituting the equation (14) into the above yields

$$
\begin{aligned}
& c_r^\sigma(x) \\
&= \lambda(m) \log u_r^{\mathrm{M}}(x,\sigma(x)) \sum_{\sigma' \in \mathcal{M}} \Bigg\{ u_r^{\mathrm{M}}(x,\sigma(x)) u_r^{\mathrm{M}}(x,\sigma'(x)) - u_r^{\mathrm{M}}(x,\sigma'(\sigma(x))) \\
&\qquad\qquad + \frac{p_r^{\mathrm{I}}(\sigma'(x))}{p_r^{\mathrm{I}}(x)} \left[u_r^{\mathrm{M}}(\sigma'(x),\sigma(x)) - u_r^{\mathrm{M}}(x,\sigma(x)) u_r^{\mathrm{M}}(\sigma'(x),x) \right] \Bigg\} \\
&+ \lambda(m) \sum_{\sigma,\sigma'} u_r^{\mathrm{M}}(x,\sigma'(x)) \left[(u_r^{\mathrm{M}} \log u_r^{\mathrm{M}} - u_r^{\mathrm{M}})(\sigma'(x),\sigma(\sigma'(x))) - (u_r^{\mathrm{M}} \log u_r^{\mathrm{M}} - u_r^{\mathrm{M}})(x,\sigma(x)) \right] \\
&= \lambda(m) \sum_{\sigma' \in \mathcal{M}} \Bigg\{ \frac{p_r^{\mathrm{I}}(\sigma'(x))}{p_r^{\mathrm{I}}(x)} \log u_r^{\mathrm{M}}(x,\sigma(x)) \left[u_r^{\mathrm{M}}(\sigma'(x),\sigma(x)) - u_r^{\mathrm{M}}(x,\sigma(x)) u_r^{\mathrm{M}}(\sigma'(x),x) \right] \\
&\qquad\qquad + u_r^{\mathrm{M}}(x,\sigma'(x)) \left[(u_r^{\mathrm{M}} \log u_r^{\mathrm{M}} - u_r^{\mathrm{M}})(\sigma'(x),\sigma(\sigma'(x))) + u_r^{\mathrm{M}}(x,\sigma(x)) \right] \\
&\qquad\qquad - u_r^{\mathrm{M}}(x,\sigma'(\sigma(x))) \log u_r^{\mathrm{M}}(x,\sigma(x)) \Bigg\}.
\end{aligned}
$$

Taking expectation and noting that $X_t^{\mathrm{M}} \overset{\text{dist}}{=} X_t^{\mathrm{I}}$, we arrive at

$$
\begin{aligned}
& \mathbb{E}\left[c_r^\sigma(X_r^{\mathrm{M}}) \right] \\
&= \lambda(m) \sum_{\sigma' \in \mathcal{M}} \sum_{x \in \mathbb{Z}_m^d} \Bigg\{ p_r^{\mathrm{I}}(\sigma'(x)) \log u_r^{\mathrm{M}}(x,\sigma(x)) \left[u_r^{\mathrm{M}}(\sigma'(x),\sigma(x)) - u_r^{\mathrm{M}}(x,\sigma(x)) u_r^{\mathrm{M}}(\sigma'(x),x) \right] \\
&\qquad\qquad + p_r^{\mathrm{I}}(x) u_r^{\mathrm{M}}(x,\sigma'(x)) \left[(u_r^{\mathrm{M}} \log u_r^{\mathrm{M}} - u_r^{\mathrm{M}})(\sigma'(x),\sigma(\sigma'(x))) + u_r^{\mathrm{M}}(x,\sigma(x)) \right] \\
&\qquad\qquad - p_r^{\mathrm{I}}(x) u_r^{\mathrm{M}}(x,\sigma'(\sigma(x))) \log u_r^{\mathrm{M}}(x,\sigma(x)) \Bigg\}.
\end{aligned}
$$

A change of variables in the first summation yields

$$
\begin{aligned}
\mathbb{E}\left[c_r^\sigma(X_r^{\mathrm{M}}) \right] &= \lambda(m) \sum_{\sigma' \in \mathcal{M}} \sum_{x \in \mathbb{Z}_m^d} p_r^{\mathrm{I}}(x) \\
&\qquad \Bigg\{ u_r^{\mathrm{M}}(x,\sigma(\sigma'(x))) \log u_r^{\mathrm{M}}(\sigma'(x),\sigma(\sigma'(x))) - u_r^{\mathrm{M}}(x,\sigma'(\sigma(x))) \log u_r^{\mathrm{M}}(x,\sigma(x)) \\
&\qquad\qquad + u_r^{\mathrm{M}}(x,\sigma'(x)) \left[u_r^{\mathrm{M}}(x,\sigma(x)) - u_r^{\mathrm{M}}(\sigma'(x),\sigma(\sigma'(x))) \right] \Bigg\} \\
&= \lambda(m) \mathbb{E}\Bigg[\sum_{\sigma' \in \mathcal{M}} \Bigg\{ u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma(\sigma'(X_r^{\mathrm{M}}))) \log u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}),\sigma(\sigma'(X_r^{\mathrm{M}}))) \\
&\qquad\qquad - u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma'(\sigma(X_r^{\mathrm{M}}))) \log u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma(X_r^{\mathrm{M}})) \\
&\qquad + u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma'(X_r^{\mathrm{M}})) \left[u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma(X_r^{\mathrm{M}})) - u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}),\sigma(\sigma'(X_r^{\mathrm{M}}))) \right] \Bigg\} \Bigg].
\end{aligned}
$$

As a result, for $0 \le s < t \le 1 - \eta$, we obtain the desired equation:

$$
\mathbb{E}\left[\varphi(u_t^{\mathrm{M}}(\sigma)) \right] - \mathbb{E}\left[\varphi(u_s^{\mathrm{M}}(\sigma)) \right] = \lambda(m) \int_s^t \sum_{\sigma' \in \mathcal{M}} \mathbb{E}\left[\mathcal{C}_r^{\mathrm{M}}(X_r^{\mathrm{M}}, u^{\mathrm{M}}; \sigma, \sigma') \right] \mathrm{d}r \,,
$$

where the operator $\mathcal{C}^{\mathrm{M}}$ defined as in Lemma 3.

B.6 Proof of main result: Theorem 2

We are now prepared to demonstrate our main convergence results. Let $\{t_k\}_{k=0}^K$ be a partition of $[0, 1-\eta]$ with mesh size $h_{k+1} = t_{k+1} - t_k$. By combining the data processing inequality [Nut21,

Lemma 1.6] with the Girsanov theorem for jump processes [CDPR25, Theorem G.0.1], we have

$$\begin{aligned}
&\mathrm{KL}(\mu_{1-\eta}|\mathrm{Law}(X^{\theta^\star}_{1-\eta}))\\
\leq&\mathrm{KL}(\mathrm{Law}(X^{\mathrm{M}}_{[0,1-\eta]})|\mathrm{Law}(X^{\theta^\star}_{[0,1-\eta]}))\\
=&\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left(q_t^{\mathrm{M}}\log\frac{q_t^{\mathrm{M}}}{q_t^{\mathrm{M},\theta^\star}}-q_t^{\mathrm{M}}+q_t^{\mathrm{M},\theta^\star}\right)(X_t^{\mathrm{M}},\sigma(X_t^{\mathrm{M}}))\right]\mathrm{d}t\,.
\end{aligned}$$

Plugging the formula of estimated generator: $q_t^{\mathrm{M},\theta^\star}(x,\sigma(x))=q_{t_k}^{\mathrm{M},\theta^\star}(X_{t_k}^{\mathrm{M}},\sigma(X_{t_k}^{\mathrm{M}}))$ for $t\in[t_k,t_{k+1})$ given $X_{t_k}^{\mathrm{M}}$ and using the shorthand $q_t^{\mathrm{M}}(\sigma)=q_t^{\mathrm{M}}(X_t^{\mathrm{M}},\sigma(X_t^{\mathrm{M}}))$ (and similarly for $q_{t_k}^{\mathrm{M}}(\sigma),q_{t_k}^{\mathrm{M},\theta^\star}(\sigma)$), we get that

$$\begin{aligned}
&\mathrm{KL}(\mu_{1-\eta}|\mathrm{Law}(X^{\theta^\star}_{1-\eta}))\\
\leq&\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left(q_t^{\mathrm{M}}\log\frac{q_t^{\mathrm{M}}}{q_{t_k}^{\mathrm{M},\theta^\star}}-q_t^{\mathrm{M}}+q_{t_k}^{\mathrm{M},\theta^\star}\right)(\sigma)\right]\mathrm{d}t\\
\leq&\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left(q_t^{\mathrm{M}}\log\frac{q_t^{\mathrm{M}}}{q_{t_k}^{\mathrm{M}}}-q_t^{\mathrm{M}}+q_{t_k}^{\mathrm{M}}+q_{t_k}^{\mathrm{M},\theta^\star}-q_{t_k}^{\mathrm{M}}+q_t^{\mathrm{M}}\log\frac{q_{t_k}^{\mathrm{M}}}{q_{t_k}^{\mathrm{M},\theta^\star}}\right)(\sigma)\right]\mathrm{d}t\,.
\end{aligned}$$

We decompose this expression into the score approximation error and a residual discretization term:

$$\begin{aligned}
&\mathrm{KL}(\mu_{1-\eta}|\mathrm{Law}(X^{\theta^\star}_{1-\eta}))\\
\leq&\underbrace{\lambda(m)\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left(\varphi(u_t^{\mathrm{M}})-\varphi(u_{t_k}^{\mathrm{M}})+\log u_{t_k}^{\mathrm{M}}(u_{t_k}^{\mathrm{M}}-u_t^{\mathrm{M}})\right)(\sigma)\right]\mathrm{d}t}_{E_{\mathrm{disc}}}\\
&+\underbrace{\sum_{k=0}^{K-1}h_{k+1}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left|q_{t_k}^{\mathrm{M},\theta^\star}(\sigma)-q_{t_k}^{\mathrm{M}}(\sigma)\right|\right]+\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\frac{q_t^{\mathrm{M}}}{q_{t_k}^{\mathrm{M}}}q_{t_k}^{\mathrm{M}}\log\frac{q_{t_k}^{\mathrm{M}}}{q_{t_k}^{\mathrm{M},\theta^\star}}(\sigma)\right]\mathrm{d}t}_{E_{\mathrm{approx}}}\,,
\end{aligned}$$

where $\lambda(m)=q(x,\sigma(x))$ denotes the fixed jump rate taking $1/2$ for NNRW and $1/m$ for URW, and the nonnegative function φ defined by $\varphi(a)=a\log(a)-a+1$ for $a\geq 0$. To handle the approximation error E_{approx}, we first note that **M**1 is indeed equivalent to

$$\sum_{k=0}^{K-1}h_{k+1}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left\{q_{t_k}^{\mathrm{M},\theta}\varphi\left(\frac{q_{t_k}^{\mathrm{M}}}{q_{t_k}^{\mathrm{M},\theta}}\right)+\left(q_{t_k}^{\mathrm{M},\theta}-q_{t_k}^{\mathrm{M}}\right)^2\right\}(X_{t_k}^{\mathrm{M}},\sigma(X_{t_k}^{\mathrm{M}}))\right]\leq\tilde{\varepsilon}^2\,,\tag{23}$$

since $X_t^{\mathrm{M}}\overset{\mathrm{dist}}{=}X_t^{\mathrm{I}}$ for all $t\in[0,1]$. Now we address the second term of E_{approx} by using the uniform bound of the score established in Lemma 1, together with the note

$$\begin{aligned}
\left|q_{t_k}^{\mathrm{M}}\log\frac{q_{t_k}^{\mathrm{M}}}{q_{t_k}^{\mathrm{M},\theta^\star}}(\sigma)\right|\leq&\left|(q_{t_k}^{\mathrm{M}}\log\frac{q_{t_k}^{\mathrm{M}}}{q_{t_k}^{\mathrm{M},\theta^\star}}-q_{t_k}^{\mathrm{M}}+q_{t_k}^{\mathrm{M},\theta^\star})(\sigma)\right|+\left|(q_{t_k}^{\mathrm{M}}-q_{t_k}^{\mathrm{M},\theta^\star})(\sigma)\right|\\
=&\,q_{t_k}^{\mathrm{M},\theta^\star}\varphi\left(\frac{q_{t_k}^{\mathrm{M}}}{q_{t_k}^{\mathrm{M},\theta^\star}}\right)(\sigma)+\left|q_{t_k}^{\mathrm{M}}(\sigma)-q_{t_k}^{\mathrm{M},\theta^\star}(\sigma)\right|\,,
\end{aligned}$$

where we used triangle inequality in the first estimate. Thereby

$$\begin{aligned}
E_{\mathrm{approx}}\lesssim&\frac{m^2}{\eta^2}\sum_{k=0}^{K-1}h_{k+1}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left|q_{t_k}^{\mathrm{M}}(\sigma)-q_{t_k}^{\mathrm{M},\theta^\star}(\sigma)\right|\right]\\
&+\frac{m^2}{\eta^2}\underbrace{\sum_{k=0}^{K-1}h_{k+1}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}q_{t_k}^{\mathrm{M},\theta^\star}\varphi\left(\frac{q_{t_k}^{\mathrm{M}}}{q_{t_k}^{\mathrm{M},\theta^\star}}\right)(\sigma)\right]}_{\leq\tilde{\varepsilon}^2\text{ by (23)}}\,.
\end{aligned}$$

We first employ Cauchy-Schwarz inequality and second apply twice the Jensen inequality for concave functions to control the remaining term of E_{approx}:

$$\begin{aligned}
&\sum_{k=0}^{K-1} h_{k+1}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left|q_{t_k}^{\mathrm{M}}(\sigma)-q_{t_k}^{\mathrm{M},\theta^\star}(\sigma)\right|\right]\\
\le&\sum_{k=0}^{K-1} h_{k+1}\mathbb{E}\left[\sqrt{|\mathcal{M}|}\sqrt{\sum_{\sigma\in\mathcal{M}}(q_{t_k}^{\mathrm{M}}(\sigma)-q_{t_k}^{\mathrm{M},\theta^\star}(\sigma))^2}\right]\\
\le&\sqrt{|\mathcal{M}|}\sum_{k=0}^{K-1} h_{k+1}\sqrt{\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}(q_{t_k}^{\mathrm{M}}(\sigma)-q_{t_k}^{\mathrm{M},\theta^\star}(\sigma))^2\right]}\\
\le&\sqrt{|\mathcal{M}|}\sqrt{\sum_{k=0}^{K-1} h_{k+1}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}(q_{t_k}^{\mathrm{M}}(\sigma)-q_{t_k}^{\mathrm{M},\theta^\star}(\sigma))^2\right]}\overset{(23)}{\le}\sqrt{|\mathcal{M}|}\tilde{\varepsilon}\,.
\end{aligned}$$

Thus we obtain the following upper bound on E_{approx}:

$$E_{\text{approx}}\lesssim m^2\eta^{-2}\sqrt{|\mathcal{M}|}\tilde{\varepsilon}+m^2\eta^{-2}\tilde{\varepsilon}^2=m^2\eta^{-2}(\sqrt{|\mathcal{M}|}\tilde{\varepsilon}+\tilde{\varepsilon}^2)\,.$$

It remains to control the time-discretization error E_{disc} to finish the proof. We write

$$\begin{aligned}
E_{\text{disc}}=\lambda(m)\underbrace{\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\log u_{t_k}^{\mathrm{M}}(u_{t_k}^{\mathrm{M}}-u_t^{\mathrm{M}})(\sigma)\right]\mathrm{d}t}_{E_1}\\
+\lambda(m)\underbrace{\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left\{\varphi(u_t^{\mathrm{M}}(\sigma))-\varphi(u_{t_k}^{\mathrm{M}}(\sigma))\right\}\right]\mathrm{d}t}_{E_2}\,,
\end{aligned}$$

To analyze the term E_2, we use the evolution of the entropy-like score established in Lemma 3:

$$\begin{aligned}
E_2&=\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\mathbb{E}\left[\sum_{\sigma\in\mathcal{M}}\left\{\varphi(u_t^{\mathrm{M}}(\sigma))-\varphi(u_{t_k}^{\mathrm{M}}(\sigma))\right\}\right]\mathrm{d}t\\
&=\lambda(m)\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\int_{t_k}^{t}\mathbb{E}\left[\sum_{\sigma,\sigma'}\left\{u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma(\sigma'(X_r^{\mathrm{M}})))\log\frac{u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}),\sigma(\sigma'(X_r^{\mathrm{M}})))}{u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma(X_r^{\mathrm{M}}))}\right.\right.\\
&\qquad\left.\left.+u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma'(X_r^{\mathrm{M}}))\left[u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma(X_r^{\mathrm{M}}))-u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}),\sigma(\sigma'(X_r^{\mathrm{M}})))\right]\right\}\right]\mathrm{d}r\mathrm{d}t\\
&\le\lambda(m)\underbrace{\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\int_{t_k}^{t_{k+1}}\mathbb{E}\left[\sum_{\sigma,\sigma'}u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma'(X_r^{\mathrm{M}}))u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma(X_r^{\mathrm{M}}))\right]\mathrm{d}r\mathrm{d}t}_{E_{2.1}}\\
&+\lambda(m)\underbrace{\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\int_{t_k}^{t_{k+1}}\mathbb{E}\left[\sum_{\sigma,\sigma'}u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma(\sigma'(X_r^{\mathrm{M}})))\left|\log\frac{u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}),\sigma(\sigma'(X_r^{\mathrm{M}})))}{u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma(X_r^{\mathrm{M}}))}\right|\right]\mathrm{d}r\mathrm{d}t}_{E_{2.2}}\,.
\end{aligned}$$

We address $E_{2.1}$ by the upper bound shown in (13):

$$\begin{aligned}
E_{2.1}&=\sum_{k=0}^{K-1}\int_{t_k}^{t_{k+1}}\int_{t_k}^{t_{k+1}}\mathbb{E}\left[\sum_{\sigma'\in\mathcal{M}}u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma'(X_r^{\mathrm{M}}))\sum_{\sigma\in\mathcal{M}}u_r^{\mathrm{M}}(X_r^{\mathrm{M}},\sigma(X_r^{\mathrm{M}}))\right]\mathrm{d}r\mathrm{d}t\\
&\lesssim\sum_{k=0}^{K-1}h_{k+1}\int_{t_k}^{t_{k+1}}\frac{d^2m^2}{(1-r)(1-t_{k+1})}\mathrm{d}r\,.
\end{aligned}$$

Let us choose $t_k = 1 - (1+h)^{-k}$ with $h > 0$, then $h_{k+1} = h(1 - t_{k+1})$, thus

$$E_{2.1} \lesssim d^2 m^2 h \int_0^{1-\eta} \frac{1}{1-r} \mathrm{d}r = d^2 m^2 h \log(\eta^{-1}) \,. \tag{24}$$

Next we invoke Lemma 1 to control the logarithmic term of $E_{2.2}$: for any $x, x' \in \mathbb{Z}_m^d$, $t, t_k \in [0, 1-\eta]$ and $\sigma \in \mathcal{M}$,

$$\left| \log \frac{u_t^{\mathrm{M}}(x, \sigma(x))}{u_{t_k}^{\mathrm{M}}(x', \sigma(x'))} \right| \lesssim \log\left(\frac{m}{1-t}\right) + \log\left(\frac{m}{1-t_k}\right) \lesssim \log\left(\frac{m}{\eta}\right) \,.$$

In addition, Lemma 1 further implies for any $x \in \mathbb{Z}_m^d$ and $r \in [t_k, t_{k+1}]$,

$$\begin{aligned}
\sum_{\sigma,\sigma'} u_r^{\mathrm{M}}(x, \sigma(\sigma'(x))) &= \sum_{\sigma,\sigma'} \mathbb{E}\left[u_r^{X_0^{\mathrm{I}} \to X_1^{\mathrm{I}}}(x, \sigma(\sigma'(x))) \middle| X_r^{\mathrm{I}} = x \right] \\
&= \mathbb{E}\left[\sum_{\sigma' \in \mathcal{M}} u_r^{X_0^{\mathrm{I}} \to X_1^{\mathrm{I}}}(x, \sigma'(x)) \sum_{\sigma \in \mathcal{M}} u_r^{X_0^{\mathrm{I}} \to X_1^{\mathrm{I}}}(\sigma'(x), \sigma(\sigma'(x))) \middle| X_r^{\mathrm{I}} = x \right] \\
&\overset{(13)}{\lesssim} \mathbb{E}\left[\frac{d^2 m^2}{(1-r)^2} \middle| X_r^{\mathrm{I}} = x \right] \leq \frac{d^2 m^2}{(1-r)(1-t_{k+1})} \,,
\end{aligned}$$

where we employed the transitivity of the bridge score in the second equality:

$$u_t^{(x_1)}(x, y) u_t^{(x_1)}(y, z) = \frac{p_{1|t}(x_1|y)}{p_{1|t}(x_1|x)} \cdot \frac{p_{1|t}(x_1|z)}{p_{1|t}(x_1|y)} = \frac{p_{1|t}(x_1|z)}{p_{1|t}(x_1|x)} = u_t^{(x_1)}(x, z) \,,$$

for any $x, y, z \in \mathbb{Z}_m^d$ and $t \in [0, 1-\eta]$. Therefore

$$\begin{aligned}
E_{2.2} &\lesssim \log(m\eta^{-1}) \sum_{k=0}^{K-1} h_{k+1} \int_{t_k}^{t_{k+1}} \frac{d^2 m^2}{(1-r)(1-t_{k+1})} \mathrm{d}r \\
&= d^2 m^2 h \log(m\eta^{-1}) \int_0^{1-\eta} \frac{1}{1-r} \mathrm{d}r = d^2 m^2 h \log(\eta^{-1}) \log(m\eta^{-1}) \,.
\end{aligned} \tag{25}$$

Combining (24) and (25) yields

$$E_2 \lesssim \lambda(m) d^2 m^2 h \log(\eta^{-1}) \log(m\eta^{-1}) \,. \tag{26}$$

To analyze E_1, we employ the evolution of the score showed in Lemma 3 to get

$$\begin{aligned}
E_1 = \lambda(m) \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} \int_{t_k}^{t} \mathbb{E}\Bigg[& \sum_{\sigma,\sigma'} \log u_{t_k}^{\mathrm{M}}(\sigma) \\
& \Big\{ u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(\sigma(X_r^{\mathrm{M}}))) - u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(X_r^{\mathrm{M}})) u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), \sigma(\sigma'(X_r^{\mathrm{M}}))) \\
& + \frac{p_r^{\mathrm{I}}(\sigma'(X_r^{\mathrm{M}}))}{p_r^{\mathrm{I}}(X_r^{\mathrm{M}})} \left[u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma(X_r^{\mathrm{M}})) u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), X_r^{\mathrm{M}}) - u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), \sigma(X_r^{\mathrm{M}})) \right] \Big\} \Bigg] \mathrm{d}r \mathrm{d}t \\
\leq \lambda(m) \sum_{k=0}^{K-1} h_{k+1} \int_{t_k}^{t_{k+1}} \mathbb{E}\Bigg[& \sum_{\sigma,\sigma'} \left| \log u_{t_k}^{\mathrm{M}}(\sigma) \right| \\
& \Big\{ \left| u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(\sigma(X_r^{\mathrm{M}}))) - u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(X_r^{\mathrm{M}})) u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), \sigma(\sigma'(X_r^{\mathrm{M}}))) \right| \\
& + \frac{p_r^{\mathrm{I}}(\sigma'(X_r^{\mathrm{M}}))}{p_r^{\mathrm{I}}(X_r^{\mathrm{M}})} \left| u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma(X_r^{\mathrm{M}})) u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), X_r^{\mathrm{M}}) - u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), \sigma(X_r^{\mathrm{M}})) \right| \Big\} \Bigg] \mathrm{d}r \,.
\end{aligned}$$

Then we employ the uniform bound of the score established in Lemma 1 to control $|\log u_{t_k}^{\mathrm{M}}(\sigma)|$:

$$
\begin{aligned}
E_1 \lesssim & \lambda(m) \log\left(m\eta^{-1}\right) \sum_{k=0}^{K-1} h_{k+1} \\
& \int_{t_k}^{t_{k+1}} \mathbb{E}\left[\sum_{\sigma,\sigma'} \left|u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(\sigma(X_r^{\mathrm{M}}))) - u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(X_r^{\mathrm{M}})) u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), \sigma(\sigma'(X_r^{\mathrm{M}})))\right|\right] \\
& \underbrace{+ \mathbb{E}\left[\sum_{\sigma,\sigma'} \frac{p_r^{\mathrm{I}}(\sigma'(X_r^{\mathrm{M}}))}{p_r^{\mathrm{I}}(X_r^{\mathrm{M}})} \left|u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma(X_r^{\mathrm{M}})) u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), X_r^{\mathrm{M}}) - u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), \sigma(X_r^{\mathrm{M}}))\right|\right]}_{E_{1.1}} \mathrm{d}r \,.
\end{aligned}
$$

The addend $E_{1.1}$ can be controlled by changing variables as follows

$$
\begin{aligned}
E_{1.1} &= \sum_{\sigma,\sigma'} \sum_{x \in \mathbb{Z}_m^d} p_r^{\mathrm{I}}(\sigma'(x)) \left|u_r^{\mathrm{M}}(\sigma'(x), \sigma(x)) - u_r^{\mathrm{M}}(x, \sigma(x)) u_r^{\mathrm{M}}(\sigma'(x), x)\right| \\
&= \sum_{\sigma,\sigma'} \sum_{x \in \mathbb{Z}_m^d} p_r^{\mathrm{I}}(x) \left|u_r^{\mathrm{M}}(x, \sigma(\sigma'(x))) - u_r^{\mathrm{M}}(\sigma'(x), \sigma(\sigma'(x))) u_r^{\mathrm{M}}(x, \sigma'(x))\right| \\
&= \sum_{\sigma,\sigma'} \mathbb{E}\left[\left|u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(X_r^{\mathrm{M}})) u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), \sigma(\sigma'(X_r^{\mathrm{M}}))) - u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(\sigma(X_r^{\mathrm{M}})))\right|\right] \,,
\end{aligned}
$$

where we leveraged the commutative property of the set of jumps $\mathcal{M}$ ($\sigma \circ \sigma' = \sigma' \circ \sigma$) in the last equality. Substituting back into E_1 leads to:

$$
\begin{aligned}
E_1 &\lesssim \lambda(m) \log\left(m\eta^{-1}\right) \sum_{k=0}^{K-1} h_{k+1} \\
&\quad \int_{t_k}^{t_{k+1}} \mathbb{E}\left[\sum_{\sigma,\sigma'} \left|u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(X_r^{\mathrm{M}})) u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), \sigma(\sigma'(X_r^{\mathrm{M}}))) - u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(\sigma(X_r^{\mathrm{M}})))\right|\right] \mathrm{d}r \\
&\lesssim \lambda(m) \log\left(m\eta^{-1}\right) \sum_{k=0}^{K-1} h_{k+1} \\
&\quad \int_{t_k}^{t_{k+1}} \mathbb{E}\left[\sum_{\sigma,\sigma'} \left\{u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(\sigma(X_r^{\mathrm{M}}))) + u_r^{\mathrm{M}}(X_r^{\mathrm{M}}, \sigma'(X_r^{\mathrm{M}})) u_r^{\mathrm{M}}(\sigma'(X_r^{\mathrm{M}}), \sigma(\sigma'(X_r^{\mathrm{M}})))\right\}\right] \mathrm{d}r \,.
\end{aligned}
$$

Arguing similarly as E_2, we obtain

$$
E_1 \lesssim \lambda(m) \log\left(m\eta^{-1}\right) \sum_{k=0}^{K-1} d^2 m^2 \frac{h_{k+1}}{1 - t_{k+1}} \int_{t_k}^{t_{k+1}} \frac{1}{1-r} \mathrm{d}r \,.
$$

Recall that $h_{k+1} = h(1 - t_{k+1})$, we arrive at

$$
\begin{aligned}
E_1 &\lesssim \lambda(m) \log\left(m\eta^{-1}\right) d^2 m^2 h \int_0^{1-\eta} \frac{1}{1-r} \mathrm{d}r \\
&= \lambda(m) d^2 m^2 h \log\left(m\eta^{-1}\right) \log\left(\eta^{-1}\right) \,. \qquad (27)
\end{aligned}
$$

Plugging the estimates for E_1 and E_2 in (26), (27) into the universal E_{disc} yields

$$
E_{\text{disc}} \lesssim (\lambda(m))^2 d^2 m^2 h \log\left(m\eta^{-1}\right) \log\left(\eta^{-1}\right) \,.
$$

Putting pieces together implies the final error bound

$$
\mathrm{KL}(\mu_{1-\eta} | \mathrm{Law}(X_{1-\eta}^{\theta^\star})) \lesssim m^2 \eta^{-2} (\sqrt{|\mathcal{M}|}\tilde{\varepsilon} + \tilde{\varepsilon}^2) + (\lambda(m))^2 d^2 m^2 h \log\left(m\eta^{-1}\right) \log\left(\eta^{-1}\right) \,,
$$

and we complete the proof of Theorem 2.

B.7 Proof of Corollary 1: Error in TV distance

We start by applying triangle inequality:

$$\|\mathrm{Law}(X_{1-\eta}^{\theta^\star}) - \mu_1\|_{\mathrm{TV}} \leq \|\mathrm{Law}(X_{1-\eta}^{\theta^\star}) - \mu_{1-\eta}\|_{\mathrm{TV}} + \|\mu_{1-\eta} - \mu_1\|_{\mathrm{TV}}\,. \tag{28}$$

The first addend can be bounded by leveraging Pinsker's inequality and Theorem 2:

$$\begin{aligned}\|\mathrm{Law}(X_{1-\eta}^{\theta^\star}) - \mu_{1-\eta}\|_{\mathrm{TV}} &\leq \sqrt{2\mathrm{KL}(\mu_{1-\eta}|\mathrm{Law}(X_{1-\eta}^{\theta^\star}))} \\ &\lesssim m\eta^{-1}\sqrt{\sqrt{|\mathcal{M}|}\tilde{\varepsilon} + \tilde{\varepsilon}^2} + \lambda(m)dm\sqrt{h\log\left(m\eta^{-1}\right)\log(\eta^{-1})}\,.\end{aligned}$$

where we used inequality $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for non-negative numbers a, b in last estimate. To conclude the proof, it remains to control the approximation error $\|\mu_{1-\eta} - \mu_1\|_{\mathrm{TV}}$ arising from the early-stopping procedure. This term represents the discrepancy between the regularized flow and the true target distribution; its analysis depends on the specific choice of base dynamics as follows.

Nearest-Neighbor Random Walk. We have

$$\begin{aligned}\|\mu_{1-\eta} - \mu_1\|_{\mathrm{TV}} &\leq \mathbb{P}\left(X_{1-\eta}^{\mathrm{I}} \neq X_1^{\mathrm{I}}\right) = 1 - \mathbb{P}(X_{1-\eta}^{\mathrm{I}} = X_1^{\mathrm{I}}) \\ &= 1 - \sum_{x_0,x_1} \pi(x_0,x_1)\frac{p_{1-\eta|0}(x_1|x_0)p_{1|1-\eta}(x_1|x_1)}{p_{1|0}(x_1|x_0)}\,.\end{aligned}$$

Thanks to the rigorous formula of transition density provided in (4), we get that

$$p_{1|1-\eta}(x_1|x_1) = \prod_{\ell=1}^{d} p^{(1)}_{1|1-\eta}(x_1^\ell|x_1^\ell) = \left(\mathrm{e}^{-\eta}\sum_{k\in\mathbb{Z}} I_{|km|}(\eta)\right)^d \geq \mathrm{e}^{-d\eta} I_0^d(\eta) \geq \mathrm{e}^{-d\eta}\,, \tag{29}$$

since

$$I_0(\eta) = \sum_{n=0}^{\infty} \frac{1}{(n!)^2}\left(\frac{\eta}{2}\right)^{2n} \geq \frac{1}{(0!)^2}\left(\frac{\eta}{2}\right)^0 = 1\,.$$

The tricky part is to manage the ratio $y_\eta = p_{1-\eta|0}(x_1|x_0)/p_{1|0}(x_1|x_0)$. By Kolmogorov backward equation and homogeneity of the base process, we have

$$\begin{aligned}y_s' &= \frac{\partial_s p_{1|s}(x_1|x_0)}{p_{1|0}(x_1|x_0)} = -\sum_{x\in\mathbb{Z}_m^d} q(x_0,x)\frac{p_{1|s}(x_1|x)}{p_{1|0}(x_1|x_0)} \\ &= d\frac{p_{1|s}(x_1|x_0)}{p_{1|0}(x_1|x_0)} - \frac{1}{2}\sum_{\sigma\in\mathcal{M}}\frac{p_{1|s}(x_1|\sigma(x_0))}{p_{1|0}(x_1|x_0)} \\ &= dy_s - \frac{1}{2}y_s\sum_{\sigma\in\mathcal{M}}\underbrace{\frac{p_{1|s}(x_1|\sigma(x_0))}{p_{1|s}(x_1|x_0)}}_{u_s^{x_0\to x_1}(x_0,\sigma(x_0))}\,.\end{aligned}$$

Thanks to the precise upper bound of the bridge score derived in the proof of Lemma 1, we get that

$$y_s' \geq dy_s - \frac{1}{2}y_s\sum_{\sigma\in\mathcal{M}}\frac{5(2m+1)}{1-s} = dy_s\left(1 - \frac{5(2m+1)}{1-s}\right)\,.$$

Therefore

$$\frac{y_s'}{y_s} \geq d\left[1 - \frac{5(2m+1)}{1-s}\right]\,.$$

Integrating from 0 to η gives

$$\log(y_\eta) - \log(y_0) \geq d\int_0^\eta \left(1 - \frac{5(2m+1)}{1-s}\right)\mathrm{d}s = d[\eta + 5(2m+1)\log(1-\eta)]\,.$$

As $y_0 = 1$, we deduce that

$$y_\eta \geq \mathrm{e}^{d\eta}(1-\eta)^{5d(2m+1)}\,. \tag{30}$$

Plugging (29) and (30) into the total variation distance yields

$$\|\mu_{1-\eta} - \mu_1\|_{\mathrm{TV}} \leq 1 - (1-\eta)^{5d(2m+1)} \sum_{x_0,x_1} \pi(x_0,x_1) = 1 - (1-\eta)^{5d(2m+1)}\,.$$

Using the Bernoulli's inequality $(1+x)^r \geq 1 + rx$ for $x \geq -1$ and $r \geq 1$, we get

$$\|\mu_{1-\eta} - \mu_1\|_{\mathrm{TV}} \leq 1 - (1 - 5d(2m+1)\eta) = 5d(2m+1)\eta \lesssim dm\eta\,.$$

Uniform Random Walk. Similarly as before,

$$\begin{aligned}\|\mu_{1-\eta} - \mu_1\|_{\mathrm{TV}} \leq& \mathbb{P}\left(X^{\mathrm{I}}_{1-\eta} \neq X^{\mathrm{I}}_1\right) = 1 - \mathbb{P}\left(X^{\mathrm{I}}_{1-\eta} = X^{\mathrm{I}}_1\right)\\ &= 1 - \sum_{x_0,x_1 \in \mathbb{Z}^d_m} \pi(x_0,x_1)\mathbb{P}(X_{1-\eta} = x_1 | X_0 = x_0, X_1 = x_1)\,.\end{aligned} \tag{31}$$

Thanks to the rigorous formula of the transition density provided in (6), we have for $x_0, x_1 \in \mathbb{Z}^d_m$:

$$\begin{aligned}&\mathbb{P}(X_{1-\eta} = x_1 | X_0 = x_0, X_1 = x_1)\\ =& \frac{p_{1-\eta|0}(x_1|x_0)p_{1|1-\eta}(x_1|x_1)}{p_{1|0}(x_1|x_0)}\\ =& \left(\frac{1-\mathrm{e}^{-(1-\eta)}}{1-\mathrm{e}^{-1}}\right)^{d_{\mathrm{H}}(x_0,x_1)} \left(\frac{1+(m-1)\mathrm{e}^{-\eta}}{m}\right)^{d} \left(\frac{1+(m-1)\mathrm{e}^{-(1-\eta)}}{1+(m-1)\mathrm{e}^{-1}}\right)^{d-d_{\mathrm{H}}(x_0,x_1)}\,.\end{aligned}$$

Since $1 + (m-1)\mathrm{e}^{-(1-\eta)} \geq 1 + (m-1)\mathrm{e}^{-1}$ and $d - d_{\mathrm{H}}(x_0,x_1) \geq 0$, the last term above larger than 1 and we deduce that

$$\mathbb{P}(X_{1-\eta} = x_1 | X_0 = x_0, X_1 = x_1) \geq \left(\frac{1-\mathrm{e}^{-(1-\eta)}}{1-\mathrm{e}^{-1}}\right)^{d_{\mathrm{H}}(x_0,x_1)} \left(\frac{1+(m-1)\mathrm{e}^{-\eta}}{m}\right)^{d}\,.$$

Now note that $\frac{1-\mathrm{e}^{-(1-\eta)}}{1-\mathrm{e}^{-1}} \leq 1$ and $d_{\mathrm{H}}(x_0,x_1) \leq d$ and apply elementary inequality $x^d \geq 1 - d(1-x)$ for $x \in [0,1]$, we obtain

$$\begin{aligned}\mathbb{P}(X_{1-\eta} = x_1 | X_0 = x_0, X_1 = x_1) &\geq \left(\frac{(1-\mathrm{e}^{-(1-\eta)})(1+(m-1)\mathrm{e}^{-\eta})}{(1-\mathrm{e}^{-1})m}\right)^{d}\\ &\geq 1 - d\left(1 - \frac{(1-\mathrm{e}^{-(1-\eta)})(1+(m-1)\mathrm{e}^{-\eta})}{(1-\mathrm{e}^{-1})m}\right)\\ &= 1 - \frac{d(1-\mathrm{e}^{-\eta})(m-1+\mathrm{e}^{-1+\eta})}{(1-\mathrm{e}^{-1})m}\\ &\geq 1 - \frac{d(1-\mathrm{e}^{-\eta})}{1-\mathrm{e}^{-1}} \geq 1 - \frac{d\eta}{1-\mathrm{e}^{-1}}\,,\end{aligned}$$

where we used inequality $\mathrm{e}^{-\eta} \geq 1 - \eta$ in the last estimate. Substituting it into (31) gives

$$\|\mu_{1-\eta} - \mu_1\|_{\mathrm{TV}} \leq 1 - \left(1 - \frac{d\eta}{1-\mathrm{e}^{-1}}\right) \sum_{x_0,x_1 \in \mathbb{Z}^d_m} \pi(x_0,x_1) = \frac{d\eta}{1-\mathrm{e}^{-1}} \lesssim d\eta\,.$$

To conclude, we plug the early-stopping error in each case into (28), yielding

$$\|\mathrm{Law}(X^{\theta^\star}_{1-\eta}) - \mu_1\|_{\mathrm{TV}} \lesssim \mathrm{C}(m)d\eta + m\eta^{-1}\sqrt{\sqrt{|\mathcal{M}|}\tilde{\varepsilon} + \tilde{\varepsilon}^2} + \lambda(m)dm\sqrt{h\log\left(m\eta^{-1}\right)\log(\eta^{-1})},$$

where the coefficient $\mathrm{C}(m)$ taking value m for NNRW and 1 for URW. We complete the proof of Corollary 1.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The claims made in the Abstract and Introduction accurately reflect the theoretical results established in the paper. We provide a detailed comparison with related work and a thorough discussion of our contributions in Section 4.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper discusses the limitations of the work performed in Section 5.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The paper provides the full set of assumptions for Theorem 2 and Corollary 1 in Section 3. Furthermore, it includes the detailed proofs of all the stated theorems in Section B and a sketch of proof of the main contributions Theorem 2 in Section 3.2.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [N/A]

Justification: Since this work is of theoretical nature, we do not provide experiments, hence our answer.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [N/A]

   Justification: Since this work is of theoretical nature, we do not provide experiments, hence our answer.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [N/A]

   Justification: Since this work is of theoretical nature, we do not provide experiments, hence our answer.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [N/A]

   Justification: Since this work is of theoretical nature, we do not provide experiments, hence our answer.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [N/A]

   Justification: Since this work is of theoretical nature, we do not provide experiments, hence our answer.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: We confirm that the research conducted our paper conform, in every respect, with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: Since this wok is of theoretical nature, it does not present societal impacts up to our knowledge.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: Sine this work is of theoretical nature, it does not present such risks.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [N/A]

    Justification: In this work, we do not use any assets to be credited.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: In the present work, we do not introduce new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: This work does not include any experiment involving human subjects, hence our answer.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The present work does not include any experiment involving human subjects, hence our answer.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLMs were used only for editing and polishing the manuscript to improve clarity and grammar. They did not impact the core methodology, scientific proofs, or original contributions of this research.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.