

A Unified Approach to Analysis and Design of Denoising Markov Models

Yinuo Ren

*Institute for Computational and Mathematical Engineering
Stanford University*

YINUOREN@STANFORD.EDU

Grant M. Rotskoff

*Department of Chemistry
Stanford University*

ROTSKOFF@STANFORD.EDU

Lexing Ying

*Department of Mathematics
Stanford University*

LEXING@STANFORD.EDU

Editor: Qiang Liu

Abstract

Probabilistic generative models based on measure transport, such as diffusion and flow-based models, are often formulated in the language of Markovian stochastic dynamics, where the choice of the underlying process impacts both algorithmic design choices and theoretical analysis. In this paper, we aim to establish a rigorous mathematical foundation for denoising Markov models, a broad class of generative models that postulate a forward process transitioning from the target distribution to a simple, easy-to-sample distribution, alongside a backward process particularly constructed to enable efficient sampling in the reverse direction. Leveraging deep connections with nonequilibrium statistical mechanics and generalized Doob's h -transform, we propose a minimal set of assumptions that ensure: (1) explicit construction of the backward generator, (2) a unified variational objective directly minimizing the measure transport discrepancy, and (3) adaptations of the classical score-matching approach across diverse dynamics. Our framework unifies existing formulations of continuous and discrete diffusion models, identifies the most general form of denoising Markov models under certain regularity assumptions on forward generators, and provides a systematic recipe for designing denoising Markov models driven by arbitrary Lévy-type processes. We illustrate the versatility and practical effectiveness of our approach through novel denoising Markov models employing geometric Brownian motion and jump processes as forward dynamics, highlighting the framework's potential flexibility and capability in modeling complex distributions.

Keywords: Generative Models, Markov Processes, Diffusion Models, Denoising Markov Models, Score-Matching, Lévy Processes

1. Introduction

Generative modeling, originating from foundational works such as VAE (Kingma, 2013) and GAN (Goodfellow et al., 2014), has become one of the central pillars of modern machine learning. Following the advent of normalizing flows (Rezende and Mohamed, 2015; Zhang et al., 2018), Markov processes emerged as a key mathematical foundation for generative

models, leading to the groundbreaking development of diffusion models (Sohl-Dickstein et al., 2015; Song and Ermon, 2019; Ho et al., 2020; Song et al., 2021b,a) and flow-based generative models (Lipman et al., 2023; Liu et al., 2023; Albergo and Vanden-Eijnden, 2023; Albergo et al., 2025). At their core, these models rely on Markovian dynamics to smoothly transform an easy-to-sample reference distribution into complex target data distributions. The adaptability of this paradigm has inspired numerous variants tailored to diverse data modalities, including discrete data (Austin et al., 2021; Campbell et al., 2022; Lou et al., 2024), and data supported on manifolds (De Bortoli et al., 2022; Huang et al., 2022; Chen and Lipman, 2024; Zhu et al., 2025). Generative models based on Markov processes have consistently achieved state-of-the-art performance across various applications, from image synthesis (Meng et al., 2022; Ho et al., 2022a; Bar-Tal et al., 2023; Chen et al., 2025c; Ma et al., 2024; Croitoru et al., 2023), video generation (Ho et al., 2022b; Yang et al., 2023; Wu et al., 2023a; Xing et al., 2024), and natural language processing (Li et al., 2022; Wu et al., 2023b; Shi et al., 2024; Xu et al., 2025), to physical, chemical, and biological domains (Xu et al., 2022; Yang et al., 2024; Frey et al., 2024; Guo et al., 2024; Stark et al., 2024).

The theoretical and algorithmic development of these models is deeply rooted in the mathematical theory of Markov processes. Continuous diffusion models, for instance, leverage the theory of stochastic differential equations (SDEs) (Tzen and Raginsky, 2019; Block et al., 2020; Chen et al., 2023b; Lee et al., 2022; Chen et al., 2023a; Benton et al., 2024a; Wang et al., 2024), along with ordinary differential equations (ODEs) through the probability flow formulation (Chen et al., 2024b; Li et al., 2024b; Gao and Zhu, 2025; Huang et al., 2025). In parallel, discrete diffusion models utilize the theory of continuous-time Markov chains and Poisson random measures (Campbell et al., 2022; Chen and Ying, 2025; Ren et al., 2025a). To further improve the efficiency and scalability of these generative frameworks, specialized numerical methods taking advantage of the properties of these underlying Markov processes have been introduced, including high-order numerical schemes (Zhang and Chen, 2023; Lu et al., 2022, 2025; Dockhorn et al., 2022; Zheng et al., 2023; Li et al., 2024a; Wu et al., 2024; Ren et al., 2025b), parallel sampling strategies (Chen et al., 2024a; Shih et al., 2024; Selvam et al., 2024), and randomized algorithms (Kandasamy and Nagaraj, 2024; Gupta et al., 2025; Li and Jiao, 2025). Recent advances in diffusion models have also significantly informed and advanced research in related guidance (Ho and Salimans, 2022; Chidambaram et al., 2024; Skreta et al., 2025; Chen et al., 2025a; Ren et al., 2026) and sampling problems (Zhang and Chen, 2022; Richter and Berner, 2024; Vargas et al., 2023; Albergo and Vanden-Eijnden, 2025; Chen et al., 2025b; Guo et al., 2025a,b).

Motivated by these remarkable empirical successes, there is a growing interest in establishing a unified theoretical understanding and generalizing current Markovian generative frameworks to accommodate a broader class of processes. Such extensions promise to facilitate the modeling of more intricate probability paths, potentially offering enhanced efficiency and expressiveness for handling complex, heavy-tailed, or multimodal distributions. Recently, Benton et al. (2024b) proposed a general framework known as *denoising Markov models*, which operates directly on the level of generators rather than explicit transition kernels to unify continuous and discrete diffusion models. Meanwhile, Holderrieth et al. (2025) developed the generator matching framework, extending flow-matching methods to general Markov processes, including pure diffusion and jump dynamics. Nevertheless, a

comprehensive generalization and theoretical characterization of denoising Markov models remain open challenges.

1.1 Contributions

This paper provides a rigorous treatment of the mathematical foundations of denoising Markov models using the language of Markov process generators. Our analysis deepens the theoretical understanding and systematically guides the algorithmic design of such generative models. Specifically, our contributions are summarized as follows:

- We employ a fundamental connection between the parametrization of backward Markov processes and the generalized Doob’s h -transform (Chetrite and Touchette, 2015). This insight allows us to formulate minimal and explicit assumptions required to justify the application of the change-of-measure argument, derive a theoretically grounded loss function that directly minimizes the KL divergence between the target and generated distributions, and generalize the classical score-matching technique (Hyvärinen and Dayan, 2005) to a broad class of Markovian models.
- Our mathematical framework unifies and extends several existing denoising Markov models, encompassing both continuous and discrete diffusion models as well as more sophisticated Lévy-Itô models (Yoon et al., 2023). Furthermore, we characterize the most general form of denoising Markov models under standard regularity conditions on their generators (Courrège, 1965) and thus generalize the permissible class of forward processes from diffusion and jump processes to arbitrary Lévy-type processes.
- Based on these theoretical developments, we provide a general yet practical recipe for training and inference in denoising Markov models. Our framework accommodates any underlying Markov process satisfying the established conditions, ensuring theoretical guarantees on model performance. Analogous to the generator matching framework (Holderrieth et al., 2025), which extends the flow-matching paradigm, our results significantly enhance the flexibility of model design and efficiency of training algorithms in the context of denoising Markov models.

In addition, we introduce novel instances of denoising Markov models by utilizing geometric Brownian motion and general jump processes as forward processes, which, to our knowledge, have not been previously explored. Through carefully designed empirical experiments, we validate the versatility, practicality, and effectiveness of our generalized framework, demonstrating its capability to model complex target distributions.

1.2 Outline

This paper is organized in the following structure: Section 2 reviews the background of denoising Markov models and recalls two key special cases, continuous and discrete diffusion models, that motivate our generalization; Section 3 develops the core mathematical theory of denoising Markov models at the generator level, the roadmap of which is illustrated in Figure 1. Section 4 specializes these general results to concrete process classes, including diffusions, finite-state jump processes, and general Lévy-type processes, illustrating how familiar score-matching losses and their discrete counterparts arise as special cases. Section 5

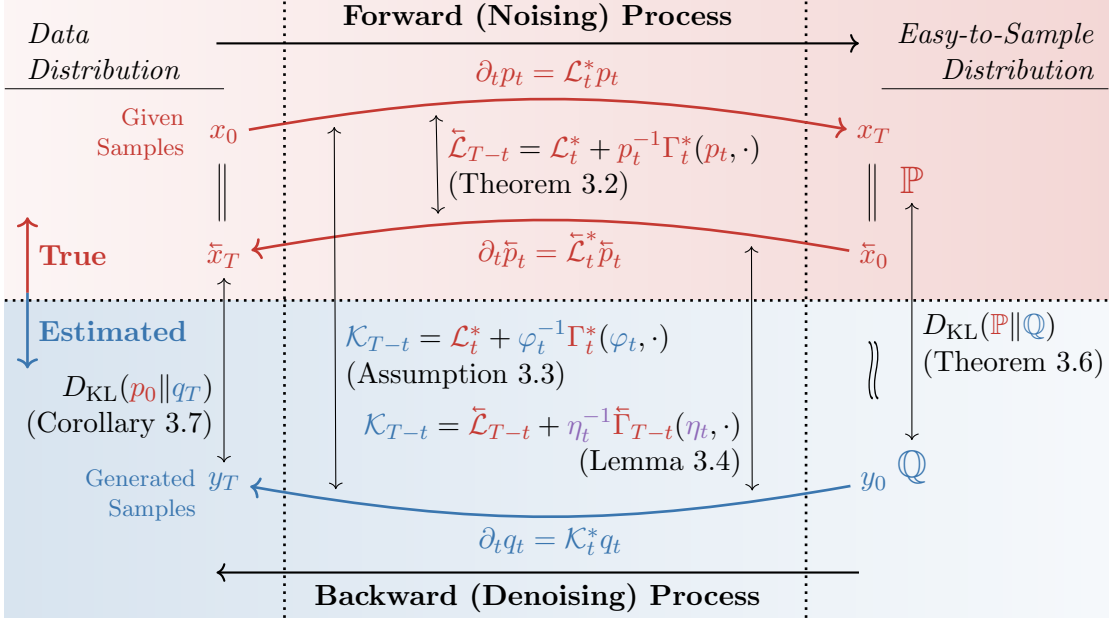


Figure 1: **Conceptual roadmap of our denoising Markov model framework.** This diagram summarizes how forward, backward, and estimated processes are related at the level of generators.

- **Red (true):** The forward Markov process x_t evolves from the data distribution p_0 to an easy-to-sample reference $p_T \approx q_0$, with generator \mathcal{L}_t and adjoint \mathcal{L}_t^* satisfying their corresponding Kolmogorov equation. The relation between the forward and backward processes is given by Theorem 3.2.
- **Blue (estimated):** The learned backward process y_t starts from q_0 and aims to reconstruct p_0 . Its generator \mathcal{K}_t is assumed to be parametrized by a positive function φ_t in the form of Assumption 3.3, where Γ_t^* is the carré du champ operator associated with the adjoint generator \mathcal{L}_t^* .
- **Purple (link):** The density ratio $\eta_t = \varphi_t p_t^{-1}$ connects the two processes via Lemma 3.4, where $\tilde{\Gamma}_t$ is the carré du champ operator associated with the backward generator $\tilde{\mathcal{L}}_t$. This demonstrates that the estimated backward generator \mathcal{K}_t can be viewed as a perturbation from the backward generator $\tilde{\mathcal{L}}_t$, which leads to the change-of-measure loss (Theorem 3.6) that upper bounds $D_{\text{KL}}(p_0 \parallel q_T)$ (Corollary 3.7).

presents proof-of-concept experiments, including models driven by geometric Brownian motion and pure jump processes, which demonstrate the flexibility of our framework beyond classical diffusion dynamics. Finally, Section 6 discusses implications and future directions.

2. Preliminaries

In this section, we outline the problem settings and basic concepts pertinent to generative modeling with Markov processes. We review the problem setting of denoising Markov models, emphasizing two prominent and well-studied diffusion model variants designed respectively for continuous and discrete data distributions.

2.1 Problem Setting and Notations

We consider the following problem setting and adopt these notational conventions throughout this paper. We refer readers to Figure 1 for a diagram of this framework. The main objective of generative modeling is to sample from a *target distribution* p_{data} supported on a measure space $(E, \mathcal{B}(E), \mu)$, given a dataset sampled from p_{data} , which forms an empirical distribution $\hat{p}_{\text{data}} \approx p_{\text{data}}$. Throughout, general states are denoted $x_t \in E$ and in bold specifically when the states have a clear multi-dimensional structure, *i.e.*, $E = \mathbb{R}^d$ or $E = [S]^d$.

In denoising Markov models, we first construct a càdlàg Markov process $(x_t)_{t \in [0, T]}$ on E in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, referred to as the *forward process* in the sequel, originating from $x_0 \sim p_{\text{data}}$. This process progressively injects noise, causing the sample to deviate from the original data distribution. We denote the distribution of x_t as p_t , with $p_0 = p_{\text{data}}$. The forward process is designed in a way that after a sufficiently long time horizon T , the resulting distribution p_T approximates or coincides with a simpler distribution q_0 that is relatively straightforward to sample from (*e.g.*, a standard Gaussian) than p_0 .

Denoising Markov models then aim to construct another denoising Markov process $(y_t)_{t \in [0, T]}$, termed the *backward process* below, whose distribution at time t is denoted as q_t . Starting from the simpler distribution q_0 , the backward process also evolves over the time horizon T to yield a distribution q_T , intended to closely approximate the target distribution p_{data} . Typically, constructing the backward process involves the following two main steps: (1) parametrizing the backward process $(y_t)_{t \in [0, T]}$ using a neural network, and (2) training the neural network by minimizing the discrepancy between the time-reversed forward process $(\tilde{x}_t)_{t \in [0, T]}$ and the backward process $(y_t)_{t \in [0, T]}$.

In general, the learnability of the backward process is significantly influenced by its parametrization, *i.e.*, the class of Markov processes representable by a particular neural network architecture. Choosing an appropriate parametrization is thus a crucial component in the design of denoising Markov models. Additionally, careful construction of the loss function is essential for enabling efficient training and providing strong theoretical guarantees. Favorably, the minimization of the loss function directly optimizes the discrepancy between the target and generated distributions, which is often quantified by the KL divergence.

In the following, we review two well-known denoising Markov models, namely continuous and discrete diffusion models, by introducing their respective forward and backward processes and the corresponding parametrization and training procedures.

2.2 (Continuous) Diffusion Models

(Continuous) diffusion models (Song et al., 2021b) are denoising Markov models designed for data supported on the d -dimensional space \mathbb{R}^d , *i.e.*, $E = \mathbb{R}^d$, $\mathcal{B}(E)$ is the Borel σ -algebra on \mathbb{R}^d , and μ is the Lebesgue measure. The forward process $(\mathbf{x}_t)_{t \in [0, T]}$ is described by a diffusion process. Specifically, the forward process $(\mathbf{x}_t)_{t \in [0, T]}$ is governed by the following stochastic differential equation (SDE):

$$d\mathbf{x}_t = \mathbf{b}_t(\mathbf{x}_t)dt + d\mathbf{w}_t, \quad (2.1)$$

where \mathbf{b}_t denotes the drift term and $(\mathbf{w}_t)_{t \geq 0}$ is a d -dimensional Wiener process. For simplicity, we assume an identity diffusion coefficient, and its generalization to the time-inhomogeneous case is straightforward with time reparametrization. A common choice of the drift is $\mathbf{b}_s = -\frac{1}{2}\mathbf{x}_s$, yielding the Ornstein-Uhlenbeck (OU) process and thus driving the forward process (2.1) towards a standard Gaussian distribution exponentially fast.

The forward SDE (2.1) naturally corresponds to a backward process $(\tilde{\mathbf{x}}_t)_{t \in [0, T]}$, which is also a diffusion process satisfying the following SDE (Anderson, 1982):

$$d\tilde{\mathbf{x}}_t = (-\mathbf{b}_{T-t}(\tilde{\mathbf{x}}_t) + \nabla \log p_{T-t}(\tilde{\mathbf{x}}_t)) dt + d\mathbf{w}_t,$$

where the term $\mathbf{s}_t(\mathbf{x}) := \nabla \log p_t(\mathbf{x})$, known as the *score function*, is often approximated by a neural network $\hat{\mathbf{s}}_t^\theta(\mathbf{x})$ parametrized by θ . In other words, the backward process is parametrized by the family of estimated backward processes $(\mathbf{y}_t)_{t \in [0, T]}$ satisfying the following SDE:

$$d\mathbf{y}_t = \left(-\mathbf{b}_{T-t}(\mathbf{y}_t) + \hat{\mathbf{s}}_{T-t}^\theta(\mathbf{y}_t) \right) dt + d\mathbf{w}_t,$$

and when the estimated score function $\hat{\mathbf{s}}_t^\theta(\mathbf{x})$ coincides with the true score function $\mathbf{s}_t(\mathbf{x})$, *i.e.*, $\hat{\mathbf{s}}_t^\theta(\mathbf{x}) = \mathbf{s}_t(\mathbf{x})$, the backward process $(\mathbf{y}_t)_{t \in [0, T]}$ is exactly the time-reversed forward process $(\tilde{\mathbf{x}}_t)_{t \in [0, T]}$.

The parameters are estimated through the score-matching objective (Hyvärinen and Dayan, 2005):

$$\min_{\theta} \mathbb{E}_{\mathbf{x}_0 \sim p_0} \left[\int_0^T \psi_t \mathbb{E}_{\mathbf{x}_t \sim p_{t|0}(\cdot|\mathbf{x}_0)} \left[\left\| \nabla \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0) - \hat{\mathbf{s}}_t^\theta(\mathbf{x}_t) \right\|^2 \right] dt \right], \quad (2.2)$$

where $p_{t|0}(\cdot|\mathbf{x}_0)$ denotes the conditional distribution at time t given a sample \mathbf{x}_0 from the data distribution p_0 , and ψ_t is a time-dependent weighting function. Specifically, in the case of the OU process where $\mathbf{b}_t = -\frac{1}{2}\mathbf{x}_t$, we have the following closed-form formula for the conditional density:

$$\nabla \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0) = -\frac{\mathbf{x}_t - \mathbf{x}_0 e^{-t/2}}{1 - e^{-t}},$$

through which Eq. (2.2) gives a practical loss function for training.

2.3 Discrete Diffusion Models

Discrete diffusion models consider data supported on a finite discrete set \mathbb{X} , *i.e.*, $E = \mathbb{X}$, $\mathcal{B}(E) = 2^{\mathbb{X}}$, and μ is the counting measure. The forward process $(x_t)_{t \in [0, T]}$ is a continuous-time Markov chain and the distribution vector $\mathbf{p}_t = (p_t(x))_{x \in \mathbb{X}}$ of the forward process at

time t satisfies the following evolution equation:

$$\frac{d\mathbf{p}_t}{dt} = \mathbf{\Lambda}_t \mathbf{p}_t, \quad (2.3)$$

where $\mathbf{\Lambda}_t = (\Lambda_t(y, x))_{x, y \in \mathbb{X}}$ is the *rate matrix* of the continuous-time Markov chain satisfying the following two conditions: (1) $\Lambda_t(x, x) = -\sum_{y \neq x} \Lambda_t(y, x)$, for all $x \in \mathbb{X}$, and (2) $\Lambda_t(x, y) \geq 0$, for all $x \neq y$. At the pointwise level, the forward process $(x_t)_{t \in [0, T]}$ also admits a stochastic integral formulation (Ren et al., 2025a):

$$x_t = x_0 + \int_0^t \int_{\mathbb{X}} (x - x_{s-}) N[\lambda](ds, dx),$$

where $N[\lambda](ds, dx)$ denotes the Poisson random measure with intensity $\lambda_s(x, x_{s-})dx$, with the intensity $\lambda_t(y, x) = \Lambda_t(y, x)(1 - \delta_x(y))$.

Strictly speaking, the jump measure depends on the outcome $\omega \in \Omega$ through the predictable state process $x_{s-}(\omega)$. A convenient and rigorous construction is via thinning of a standard Poisson random measure on an augmented space. Let μ be a reference measure on \mathbb{X} (e.g., counting measure when \mathbb{X} is discrete), and let $N(ds, dx, dz)$ be a Poisson random measure on $[0, \infty) \times \mathbb{X} \times \mathbb{R}_+$ with deterministic intensity $ds\mu(dx)dz$, meaning that (1) for any $0 \leq s < t < \infty$ and any Borel sets $A \subset \mathbb{X}$ and $B \subset \mathbb{R}_+$ with finite measure,

$$N((s, t] \times A \times B) \sim \text{Poisson}((t - s)\mu(A)|B|),$$

and (2) the random variables $N((s, t] \times A_i \times B_i)$ are independent over disjoint sets $(s, t] \times A_i \times B_i$. Given a predictable nonnegative intensity $\lambda_s(x, x_{s-}(\omega))$, define the thinned (integer-valued) random measure on $[0, \infty) \times \mathbb{X}$ by

$$N[\lambda](\omega, ds, dx) := \int_0^\infty \mathbf{1}_{z \leq \lambda_s(x, x_{s-}(\omega))} N(ds, dx, dz).$$

Then $N[\lambda]$ has predictable compensator $\lambda_s(x, x_{s-}(\omega))ds\mu(dx)$. We refer readers to Böttcher et al. (2013) for standard Poisson random measures and to Protter (1983); Ren et al. (2025a) for more details on this construction procedure and applications. In the following we omit the dependence on ω for notational simplicity.

It is a classical result that the forward process (2.3) also corresponds to a backward process $(\tilde{x}_t)_{t \in [0, T]}$ with another rate matrix $\bar{\mathbf{\Lambda}}_t = (\bar{\Lambda}_t(y, x))_{x, y \in \mathbb{X}}$ defined as

$$\bar{\Lambda}_t(y, x) = \begin{cases} s_{T-t}(x, y)\Lambda_{T-t}(x, y), & \forall x \neq y, \\ -\sum_{y' \neq x} \bar{\Lambda}_t(y', x), & \forall x = y, \end{cases} \quad (2.4)$$

where the score function \mathbf{s}_t , a vector-valued function, is defined as

$$\mathbf{s}_t(x) = (s_t(x, y))_{y \in \mathbb{X}} = \frac{\mathbf{p}_t}{p_t(x)}.$$

Similar to the continuous case, we use a neural network $\widehat{\mathbf{s}}_t^\theta$ with parameters θ to estimate the score function \mathbf{s}_t , i.e., parametrizing the backward process $(\tilde{x}_t)_{t \in [0, T]}$ by the family of estimated backward processes $(y_t)_{t \in [0, T]}$ satisfying the following evolution equation:

$$\frac{dq_t}{dt} = \bar{\mathbf{\Lambda}}_t \mathbf{q}_t, \quad \text{with} \quad \bar{\Lambda}_t(y, x) = \begin{cases} \widehat{\mathbf{s}}_{T-t}^\theta(x, y)\Lambda_{T-t}(x, y), & \forall x \neq y, \\ -\sum_{y' \neq x} \bar{\Lambda}_t(y', x), & \forall x = y. \end{cases}$$

This also corresponds to a pointwise stochastic integral formulation

$$y_t = y_0 + \int_0^t \int_{\mathbb{X}} (y - y_{s-}) N[\bar{\lambda}](ds, dy),$$

where $N[\bar{\lambda}](ds, dy)$ is the Poisson random measure with evolving intensity $\bar{\lambda}_s(y, y_{s-})dy$, with the intensity $\bar{\lambda}_t(y, x) = \bar{\Lambda}_t(y, x)(1 - \delta_x(y))$.

The score function \mathbf{s}_t is estimated via the score-matching objective (Lou et al., 2024):

$$\min_{\theta} \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T \psi_t \mathbb{E}_{x_t \sim p_{t|0}(\cdot|x_0)} \left[\sum_{y \neq x_t} \left(\widehat{\mathbf{s}}_t^{\theta}(x_t, y) - \frac{p_{t|0}(y|x_0)}{p_{t|0}(x_t|x_0)} \log \widehat{\mathbf{s}}_t^{\theta}(x_t, y) \right) \Lambda_t(x_t, y) \right] \right] dt. \quad (2.5)$$

In practice, a usual setting of the state space is the set of d -dimensional discrete vectors, *i.e.*, $\mathbb{X} = [S]^d$, where S is the number of possible states along each dimension. The forward process $(\mathbf{x}_t)_{t \in [0, T]}$ is often chosen such that the rate matrix Λ_t is sparse, for which two rate matrices are commonly adopted: (1) the uniform rate matrix: $\lambda(\mathbf{y}, \mathbf{x}) = d^{-1}$ only if $\|\mathbf{y} - \mathbf{x}\|_0 = 1$, for which we have

$$p_{t|0}(\mathbf{x}_t | \mathbf{x}_0) = \prod_{i=1}^d \left(e^{-t} \delta_{x_0^i}(x_t^i) + \frac{1 - e^{-t}}{S} \right),$$

and (2) the masked rate matrix: $\lambda(\mathbf{x}_{-k} \oplus 0, \mathbf{x}) = 1$ if $x_k \neq 0$, where 0 is appended to each dimension of the state space as a null state and $\mathbf{x}_{-k} \oplus 0$ denotes the state x with the k -th dimension set to 0, in which case we have

$$p_{t|0}(\mathbf{x}_t | \mathbf{x}_0) = \prod_{i=1}^d \left(e^{-t} \delta_{x_0^i}(x_t^i) + (1 - e^{-t}) \delta_0(x_t^i) \right).$$

Both cases converge exponentially fast, and Eq. (2.5) gives a practical loss function for training.

3. Denoising Markov Models

In this section, we present the mathematical framework underpinning denoising Markov models. We first introduce the concept of generators, which serve as fundamental building blocks of these models. Next, we discuss the time reversal of the forward process, yielding the backward process, explore how its generator relates explicitly to the generator of the forward process, and provide appropriate assumptions regarding the parametrization of the estimated backward process. Connecting the relations between the forward and backward processes to the generalized Doob's h -transform, we provide the change-of-measure arguments that naturally lead to a derivation of the score-matching objective for general denoising Markov models formulated in terms of generators. Finally, we give meta-algorithms for training and sampling from denoising Markov models, which can be applied to any forward process satisfying the established conditions, and provide a theoretical guarantee on the performance of the estimated backward process in terms of the KL divergence between the target and generated distributions.

Throughout this section, definitions, theorems, and proofs are presented informally to highlight the essential concepts and intuitions. For precise mathematical definitions, rigorous theorem statements, explicit assumptions, and detailed proofs, we direct the reader to Appendix A.

3.1 Mathematical Background

We assume that the forward process $(x_t)_{t \in [0, T]}$ is governed by an *evolution system* $(U_{t,s})_{s \leq t}$ (Definition A.1), defined by

$$U_{t,s}f(x) = \mathbb{E}[f(x_t)|x_s = x].$$

Furthermore, we assume that $(U_{t,s})_{s \leq t}$ is a *Feller evolution system* (Definition A.2). Roughly speaking, a Feller evolution system is a family of linear, time-evolution operators mapping the space $C_0(E)$ consisting of all continuous functions vanishing at infinity on E , into itself, that satisfy positivity-preservation, contractivity, and strong continuity.

Feller evolution systems can be characterized by their right generators. The right generator \mathcal{L}_t of the forward process $(x_t)_{t \in [0, T]}$ is defined as follows:

Definition 3.1 (Forward Generator, Informal Version of Definition A.3) *For each $t \in [0, T]$, the right generator of the forward process $(x_t)_{t \in [0, T]}$ at time t is given by:*

$$\mathcal{L}_t f(x) = \lim_{h \rightarrow 0^+} \mathbb{E} \left[\frac{f(x_{t+h}) - f(x_t)}{h} \middle| x_t = x \right].$$

Should no ambiguity arise, we refer simply to this as the forward generator. If the forward process (x_t) is time-homogeneous, then $T_{t-s} := U_{t,s}$ forms a *one-parameter semigroup* or simply semigroup (Definition A.4), with a time-homogeneous generator \mathcal{L} (Definition A.5). In the special case when $(U_{t,s})_{s \leq t}$ is a Feller evolution system, the corresponding semigroup $(T_t)_{t \geq 0}$ is called a *Feller semigroup* (Definition A.6).

Feller processes, which are Markov processes governed by Feller semigroups (Definition A.7), have been extensively studied in the literature (Ethier and Kurtz, 2009; Böttcher et al., 2013). They possess several convenient properties, including the existence of transition kernels as unique positive Radon measures by the Riesz representation theorem (Böttcher et al., 2013, Theorem 1.5), an adapted form of the Hille-Yosida theorem (Ethier and Kurtz, 2009, Theorem 4.2.2), and the following Dynkin's formula (Theorem A.8):

$$\mathbb{E}[f(x_t)|x_0 = x] - f(x_0) = \mathbb{E} \left[\int_0^t \mathcal{L}f(x_s) ds \middle| x_0 = x \right].$$

In most practical settings, however, the evolution system $(U_{t,s})_{s \leq t}$ is time-inhomogeneous. Nevertheless, one can often transform the forward process $(x_t)_{t \in [0, T]}$ into an augmented process that incorporates the time index explicitly, without introducing additional randomness. We refer to Definition A.9 for details of this transformation. By Theorem A.14, the augmented process $(\tilde{x}_t)_{t \in [0, T]}$, defined on the probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ with the augmented generator $\tilde{\mathcal{L}}$, is a Feller process if and only if the original process $(x_t)_{t \in [0, T]}$ is governed by a

Feller evolution system. Under certain smoothness conditions, the generators of the original and augmented processes are connected as follows (Proposition A.12):

$$\tilde{\mathcal{L}}f(t, x) = \lim_{h \rightarrow 0^+} \left[\frac{f(t+h, x_{t+h}) - f(t, x_t)}{h} \Big|_{x_t = x} \right] = \partial_t f(t, x) + \mathcal{L}_t f(t, x).$$

Assuming the existence of the adjoint of the forward generator \mathcal{L}_t (Definition A.19), denoted as \mathcal{L}_t^* , we recall the classical Kolmogorov forward equation (Theorem A.23):

$$\partial_t p_t = \mathcal{L}_t^* p_t.$$

3.2 Time-Reversal of Forward Process

As discussed in Sections 2.2 and 2.3, once the forward process $(x_t)_{t \in [0, T]}$ has been constructed with the marginal distributions satisfying $p_0 = p_{\text{data}}$ and $p_T \approx q_0$, the objective is to construct a corresponding backward process $(y_t)_{t \in [0, T]}$, whose marginal distribution q_T at time T is close to p_0 .

Due to the Markov property of the forward process $(x_t)_{t \in [0, T]}$, there exists a natural *time-reversal process* $(\tilde{x}_t)_{t \in [0, T]}$ (Definition A.17), which we refer to below as the (true) backward process. Specifically, for each $t \in [0, T)$, we define $\tilde{x}_t = x_{(T-t)^-}$, where x_{t^-} denotes the left limit at time t . The time-reversal process thus has the initial distribution $\tilde{x}_0 \sim p_T$ and terminal distribution $\tilde{x}_T := x_0 \sim p_0$. It can be verified that the time-reversal process $(\tilde{x}_t)_{t \in [0, T]}$ of a càdlàg Markov process $(x_t)_{t \in [0, T]}$ remains càdlàg and Markovian, thereby admitting a generator called the backward generator, denoted by $\tilde{\mathcal{L}}_t$, defined analogously to the forward generator \mathcal{L}_t in Definition 3.1.

The relationship between the forward generator \mathcal{L}_t and the backward generator $\tilde{\mathcal{L}}_t$ is characterized by the following result:

Theorem 3.2 (Time-Reversal of Forward Process (Cattiaux et al., 2023), Informal Version of Theorem A.22) *Under certain regularity assumptions on the densities p_t (Assumption A.21), the backward generator $\tilde{\mathcal{L}}_t$ can be explicitly expressed in terms of the forward generator \mathcal{L}_t and the density p_t at time t as follows:*

$$p_t \tilde{\mathcal{L}}_{T-t} f = p_t \mathcal{L}_t^* f + \Gamma_t^*(p_t, f), \quad (3.1)$$

where \mathcal{L}_t^* is the adjoint operator of \mathcal{L}_t (Definition A.19), and

$$\Gamma_t^*(p_t, f) := \mathcal{L}_t^*(p_t f) - p_t \mathcal{L}_t^* f - f \mathcal{L}_t^* p_t$$

is the carré du champ operator associated with the adjoint operator \mathcal{L}_t^* .

Although Theorem 3.2 provides an explicit formula connecting the backward generator $\tilde{\mathcal{L}}_t$ to the forward generator \mathcal{L}_t , the expression inherently involves the unknown density p_t at each time t . To overcome this practical limitation, we introduce another Markov process $(y_t)_{t \in [0, T]}$, driven by another generator \mathcal{K}_t . Typically, we restrict the possible choices of the generator \mathcal{K}_t to a parametrized class of generators that is sufficiently flexible to contain a broad spectrum of generators, especially including the true backward generator $\tilde{\mathcal{L}}_t$, while remaining computationally tractable and learnable from available data samples from p_0 and simulations of the forward generator \mathcal{L}_t .

Motivated by the particular form involving the carré du champ operator associated with the adjoint generator \mathcal{L}_t^* in the backward generator $\tilde{\mathcal{L}}_{T-t}$ given by Eq. (3.1), the following assumption constitutes a central structural hypothesis in the denoising Markov model framework:

Assumption 3.3 (Parametrization of Backward Generator, Informal Version of Assumption A.24) *Under certain regularity conditions on η_t (Assumption A.25), we assume the estimated backward generator \mathcal{K}_t , associated with the estimated backward process $(y_t)_{t \in [0, T]}$, is parametrized by a strictly positive function $\varphi_t : E \rightarrow \mathbb{R}^+$ satisfying*

$$\varphi_t \mathcal{K}_{T-t} f = \varphi_t \mathcal{L}_t^* f + \Gamma_t^*(\varphi_t, f). \quad (3.2)$$

In practical implementations, we parametrize either φ_t itself or its equivalent coordinates by neural networks, whose parameters are denoted by θ . The corresponding parametrized generator is denoted by \mathcal{K}_t^θ .

Observe that when the function φ_t coincides exactly with the marginal density p_t , the estimated backward generator \mathcal{K}_t reduces to the true backward generator $\tilde{\mathcal{L}}_t$. Therefore, the strictly positive function φ_t intrinsically represents an approximation to the marginal densities. More generally, defining the density ratio of φ_t and p_t by $\eta_t = \varphi_t p_t^{-1}$, and assuming certain regularity conditions thereof (Assumption A.25), we can relate the estimated backward generator \mathcal{K}_t to the true backward generator $\tilde{\mathcal{L}}_{T-t}$ through the following lemma:

Lemma 3.4 (Informal Version of Lemma A.26) *The estimated backward generator \mathcal{K}_t and the true backward generator $\tilde{\mathcal{L}}_t$ are related through the following relation:*

$$\eta_t \mathcal{K}_{T-t} f = \eta_t \tilde{\mathcal{L}}_{T-t} f + \tilde{\Gamma}_{T-t}(\eta_t, f),$$

where $\tilde{\Gamma}_t(f, g) = \tilde{\mathcal{L}}_t(fg) - f\tilde{\mathcal{L}}_t g - g\tilde{\mathcal{L}}_t f$ is the carré du champ operator associated with the time-reversal generator $\tilde{\mathcal{L}}_t$.

This lemma offers a convenient way of understanding and analyzing the parametrized backward generator \mathcal{K}_t relative to the unknown true backward dynamics.

3.3 Change of Measure

The relation between the estimated backward generator \mathcal{K}_t and the true backward generator $\tilde{\mathcal{L}}_t$, outlined in Lemma 3.4, can be viewed as applying a ‘‘perturbation’’ of the form $\eta_{T-t}^{-1} \tilde{\Gamma}_t(\eta_{T-t}, \cdot)$ to $\tilde{\mathcal{L}}_t$. This particular perturbation, involving the carré du champ operator, repeatedly appears in the connection between the backward generator $\tilde{\mathcal{L}}_t$ and the adjoint operator \mathcal{L}_t^* (Theorem 3.2), as well as in our parametrization for the estimated backward generator \mathcal{K}_t (Assumption 3.3). These relationships are summarized in Figure 1. The question remains as to how this specific form provides a quantitative measure of perturbation, enabling the derivation of practical loss functions and error bounds for denoising Markov models.

As explored in the literature (Kunita, 1969; Ethier and Kurtz, 2009; Fleming and Soner, 2006; Palmowski and Rolski, 2002), this form of perturbation closely relates to the *generalized Doob’s h-transform* (Chetrite and Touchette, 2015). The main idea of the transform is summarized in the following theorem:

Theorem 3.5 (Generalized Doob's h -Transform (Chetrite and Touchette, 2015))

Suppose \mathcal{L}_t is the generator of a Markov process $(x_t)_{t \in [0, T]}$ with path measure \mathbb{P} , h_t is a strictly positive and sufficiently smooth function, and λ_t is an arbitrary function, for any $t \in [0, T]$. Then the transformed process with the generator $\mathcal{L}_t^{h, \lambda}$, defined as

$$\mathcal{L}_t^{h, \lambda} f = h_t^{-1} \mathcal{L}_t(h_t f) - \lambda_t f,$$

is another Markov process with path measure $\mathbb{P}^{h, \lambda}$ absolutely continuous with respect to \mathbb{P} , with the Radon-Nikodym derivative satisfying

$$\frac{d\mathbb{P}^{h, \lambda}}{d\mathbb{P}}(x_{[0, T]}) = \frac{h_T(x_T)}{h_0(x_0)} \exp\left(-\int_0^T (h_t^{-1} \partial_t h_t + \lambda_t)(x_t) dt\right), \quad (3.3)$$

where $x_{[0, T]}$ denotes the path of the process $(x_t)_{t \in [0, T]}$.

In particular, choosing $\lambda_t = h_t^{-1} \mathcal{L}_t h_t$ results in a conservative transformed process, i.e., $\mathcal{L}_t^{h, \lambda} 1 = 0$, satisfying

$$\mathcal{L}_t^{h, \lambda} f = h_t^{-1} \mathcal{L}_t(h_t f) - h_t^{-1} f \mathcal{L}_t h_t = \mathcal{L}_t f + h_t^{-1} \Gamma_t(h_t, f),$$

where $\Gamma_t(h_t, f) = \mathcal{L}_t(h_t f) - h_t \mathcal{L}_t f - f \mathcal{L}_t h_t$ is the carré du champ operator associated with the generator \mathcal{L}_t .

The generalized Doob's transform is broadly used in physical contexts, as it is related to stochastic mechanics (Meyer and Zheng, 1985) and fluctuation-dissipation relations (Chetrite and Gupta, 2011). It is also explored in the study of the Metropolis algorithm (Diaconis and Miclo, 2009). For non-perturbative generalization, we refer to Jarzynski (1997); Gallavotti and Cohen (1995).

In the denoising Markov model context, Lemma 3.4 can be rewritten, after replacing t by $T - t$, as

$$\mathcal{K}_t f = \eta_{T-t}^{-1} \tilde{\mathcal{L}}_t(\eta_{T-t} f) - \eta_{T-t}^{-1} f \tilde{\mathcal{L}}_t \eta_{T-t}.$$

Therefore, \mathcal{K}_t is precisely the generalized Doob's h -transform of the true backward generator $\tilde{\mathcal{L}}_t$ with

$$h_t = \eta_{T-t}, \quad \lambda_t = \eta_{T-t}^{-1} \tilde{\mathcal{L}}_t \eta_{T-t}.$$

In particular, the approximate backward process may be viewed as a path-space reweighting of the true backward process, with the density ratio η_t acting as the reweighting factor. When this reweighting is trivial (i.e., $\eta_t = 1$), the estimated backward generator \mathcal{K}_t coincides with the true backward generator $\tilde{\mathcal{L}}_t$.

This observation leads directly to the following change-of-measure identity, which provides a quantitative measure of the perturbation between the estimated backward generator \mathcal{K}_t and the true backward generator $\tilde{\mathcal{L}}_t$ in terms of a variational form of the density ratio η_t . The perturbation should originally involve $\tilde{\mathcal{L}}_t$ (cf. Theorem A.29) as in Eq. (3.3) when perturbing \mathcal{L} but thanks to Theorem 3.2, it can be rewritten in terms of the forward generator \mathcal{L}_t as follows:

Theorem 3.6 (Change of Measure, Informal Version of Corollary A.30) *Under certain regularity conditions on η_t (Assumption A.25), there exists a probability measure \mathbb{Q}*

absolutely continuous to the original probability measure \mathbb{P} with the Radon-Nikodym derivative satisfying that

$$D_{\text{KL}}(\mathbb{P}||\mathbb{Q}) = \mathbb{E}_{\mathbb{P}} \left[\log \frac{d\mathbb{P}}{d\mathbb{Q}} \right] = \mathbb{E}_{\mathbb{P}} \left[\int_0^T (\eta_t \mathcal{L}_t \eta_t^{-1} + \mathcal{L}_t \log \eta_t)(x_t) dt \right] := \mathfrak{L}[\eta_t], \quad (3.4)$$

under which the time-reversal process $(\tilde{x}_t)_{t \in [0, T]}$ is governed by the generator \mathcal{K}_t , where the expectation is taken over the forward process $(x_t)_{t \in [0, T]}$.

We will omit the subscript \mathbb{P} indicating the expectation over the forward process $(x_t)_{t \in [0, T]}$ in the following discussion for brevity. The result in Theorem 3.6 directly yields error bounds for generative models in terms of KL divergence, as shown by the following corollary by applying the data-processing inequality and the chain rule of KL divergence:

Corollary 3.7 *The following error bound holds for the denoising Markov model:*

$$D_{\text{KL}}(p_0||q_T) \leq D_{\text{KL}}(p_T||q_0) + \mathfrak{L}[\eta_t],$$

where $\mathfrak{L}[\eta_t]$ is defined as in Eq. (3.4).

This result highlights that the quantity $\mathfrak{L}[\eta_t]$ not only quantifies the perturbation between the estimated backward generator \mathcal{K}_t and the true backward generator $\tilde{\mathcal{L}}_t$ but also naturally serves as an appropriate loss function for training denoising Markov models. By minimizing this loss, we directly obtain an error bound on the generative model, measured by the KL divergence between the target distribution p_0 and the estimated distribution q_T , up to an additional term $D_{\text{KL}}(p_T||q_0)$ that typically decays exponentially and reflects the convergence of the forward process.

To facilitate practical and efficient training, the marginal densities $p_t(x_t)$ can be expressed instead in terms of the conditional distributions $p_{t|0}(x_t|x_0)$, which are often simpler to compute or available in closed-form solutions. This substitution leads to a score-matching objective tailored for general denoising Markov models, as captured in the following corollary, extending the classical score-matching framework (Hyvärinen and Dayan, 2005; Vincent, 2011).

Corollary 3.8 (Score-Matching) *Define the conditional density ratio as*

$$\eta_{t|0}(\cdot|x_0) = \varphi_t(\cdot) p_{t|0}^{-1}(\cdot|x_0),$$

where $p_{t|0}(\cdot|x_0)$ denotes the conditional distribution of x_t given x_0 . Then, the score-matching loss is given by

$$\mathfrak{L}_{\text{SM}}[\eta_{t|0}] = \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T \mathbb{E}_{x_t \sim p_{t|0}(\cdot|x_0)} \left[\left(\eta_{t|0}(\cdot|x_0) \mathcal{L}_t \eta_{t|0}^{-1}(\cdot|x_0) + \mathcal{L}_t \log \eta_{t|0}(\cdot|x_0) \right) (x_t) \right] dt \right], \quad (3.5)$$

and when the function η_t depends on the parameters θ through the parametrization of φ_t by its neural network representation, the score-matching objective satisfies that

$$\arg \min_{\theta} D_{\text{KL}}(\mathbb{P}||\mathbb{Q}^{\theta}) = \arg \min_{\theta} \mathfrak{L}[\eta_t^{\theta}] = \arg \min_{\theta} \mathfrak{L}_{\text{SM}}[\eta_{t|0}^{\theta}],$$

where constants independent of the model parameters θ are omitted.

In general, due to the non-negativity of the integrands of $\mathfrak{L}[\eta_t]$ (3.4) and $\mathfrak{L}_{\text{SM}}[\eta_{t|0}]$ (3.5) (cf. (Benton et al., 2024b, Proposition 1(a))), one may reweight the integrands by a function ψ_t to mitigate the concentration of the loss in regions where its magnitude is large (Song et al., 2021b), similar to the approach in Eqs. (2.2) and (2.5). Practically, this reweighting can be implemented by sampling time t from a non-uniform distribution Ψ over $[0, T]$ during training (cf. Algorithm 1), resulting in the following time-reweighted loss:

$$\mathfrak{L}_{\text{SM}, \Psi}[\eta_{t|0}] = \mathbb{E}_{x_0 \sim p_0, t \sim \Psi} \left[\mathbb{E}_{x_t \sim p_{t|0}(\cdot|x_0)} \left[\left(\eta_{t|0}(\cdot|x_0) \mathcal{L}_t \eta_{t|0}^{-1}(\cdot|x_0) + \mathcal{L}_t \log \eta_{t|0}(\cdot|x_0) \right) (x_t) \right] \right].$$

Remark 3.9 (Comparison with Benton et al. (2024b)) *We briefly compare our proposed framework and the approach presented by Benton et al. (2024b). While our framework relies on a single explicit assumption (Assumption 3.3), Benton et al. (2024b) propose two independent assumptions concerning the parametrization of the estimated backward generator \mathcal{K}_t , which, in our notation, can be expressed as follows:*

- (1) *For each $t \in [0, T]$, there exists an auxiliary Markov process $(z_t)_{t \in [0, T]}$ governed by a Feller evolution system with generator \mathcal{M}_t , together with a function c_t , such that*

$$\mathcal{K}_{T-t}^* f = \mathcal{M}_{T-t} f + c_{T-t} f$$

for all sufficiently regular f ;

- (2) *For each $t \in [0, T]$, there exists a strictly positive function β_t satisfying*

$$\beta_t^{-1} \mathcal{M}_{T-t} f = \mathcal{L}_t(\beta_t^{-1} f) - f \mathcal{L}_t \beta_t^{-1}.$$

Under these assumptions, Benton et al. (2024b) apply a change-of-measure argument between the original process $(x_t)_{t \in [0, T]}$ and the auxiliary process $(z_t)_{t \in [0, T]}$, and subsequently invoke the integral-by-parts formula for semimartingales and the Feynman-Kac formula to handle the additional term c_t .

In contrast, our simpler and more explicit assumption (Assumption 3.3) inherently satisfies the assumptions above. Specifically, using Eq. (3.2) and the definition of the adjoint operator, we have

$$\mathcal{K}_{T-t}^* g = \varphi_t \mathcal{L}_t(\varphi_t^{-1} g) - \varphi_t^{-1} g \mathcal{L}_t^* \varphi_t = \varphi_t \mathcal{L}_t(\varphi_t^{-1} g) - \varphi_t g \mathcal{L}_t \varphi_t^{-1} + g (\varphi_t \mathcal{L}_t \varphi_t^{-1} - \varphi_t^{-1} \mathcal{L}_t^* \varphi_t),$$

which corresponds precisely to their assumptions with the choices

$$\beta_t = \varphi_t, \quad \mathcal{M}_{T-t} g = \varphi_t \mathcal{L}_t(\varphi_t^{-1} g) - \varphi_t g \mathcal{L}_t \varphi_t^{-1}, \quad \text{and} \quad c_{T-t} = \varphi_t \mathcal{L}_t \varphi_t^{-1} - \varphi_t^{-1} \mathcal{L}_t^* \varphi_t.$$

Moreover, while the choice of c_t is left implicit in the approach of Benton et al. (2024b), we explicitly characterize c_t in our framework and demonstrate via Lemma 3.4 and Theorem 3.5 that this selection is uniquely determined to ensure the conservativeness of the estimated backward process $(y_t)_{t \in [0, T]}$. By clearly specifying c_t , our framework provides a direct and explicit representation of the estimated backward generator \mathcal{K}_t , thereby enabling a broader class of denoising Markov models with arbitrary Feller forward processes, as elaborated in Section 4 and verified by several examples in Section 5.

Furthermore, our analysis reveals explicit and intuitive relationships between the forward generator \mathcal{L}_t , the true backward generator $\tilde{\mathcal{L}}_t$, and the estimated backward generator \mathcal{K}_t , as illustrated clearly in Figure 1. By interpreting these relationships through a generalized Doob’s h -transform (Theorem 3.2, Assumption 3.3, and Lemma 3.4), our framework avoids the necessity of introducing an auxiliary process $(z_t)_{t \in [0, T]}$. Instead, we directly employ a streamlined change-of-measure argument (Theorem 3.6) between the true backward process $(\tilde{x}_t)_{t \in [0, T]}$ and its estimate $(y_t)_{t \in [0, T]}$. This simplification significantly clarifies the analysis, enhances the conceptual transparency of the method, and provides clearer insights for practical implementation in designing denoising Markov models.

3.4 Meta-Algorithm

We now introduce the meta-algorithm for denoising Markov models and outline several key principles underlying their design.

Training Process. The training process (Algorithm 1) of denoising Markov models typically relies on four crucial components: (1) an empirical distribution \hat{p}_0 approximating the target distribution p_0 , (2) a simple, easy-to-sample reference distribution to which the forward process rapidly converges after time horizon T , (3) an explicit formula for the conditional distribution $p_{t|0}(\cdot|x_0)$, and (4) an efficient method for simulating the forward process $(x_t)_{t \in [0, T]}$ generated by \mathcal{L}_t .

In generative modeling scenarios, the first requirement is naturally fulfilled by available data samples. However, the remaining conditions depend on careful choices of the forward process. Choosing the transition kernel $p_{t|0}$ to be space-homogeneous and even further time-homogeneous is common in practice to simplify its computation. We provide concrete examples addressing these criteria in Section 4.

Inference Process. Inference for denoising Markov models is straightforward in principle: given the explicit form of the estimated backward generator \mathcal{K}_t^θ (specific examples again appear in Section 4), we simulate the estimated backward process $(y_t)_{t \in [0, T]}$. Starting from the initial distribution q_0 , samples are generated by numerically discretizing and simulating this continuous-time process.

Since practical implementations involve discretization, controlling the resulting numerical error becomes crucial for the model’s performance. To quantify the numerical error, we consider a uniform discretization of the interval $[0, T]$ into time steps of length κ , assuming $T = L\kappa$ for some integer L . For $\ell \in [L]$, let $\hat{q}_{\ell\kappa}$ denote the distribution of samples after ℓ discretization steps. The inference procedure for the discretized denoising Markov model is detailed in Algorithm 2. We impose the following assumption on the numerical error introduced at each discretization step:

In the following theorem, we consider a general case where the time horizon T is discretized into time steps of equal length κ . For simplicity, we assume that $T = L\kappa$ for some integer L . For any $\ell \in [L]$, we denote the distribution of obtained samples after ℓ steps of simulation by $\hat{q}_{\ell\kappa}$. We refer to Algorithm 2 for the inference procedure of the denoising Markov model. Within each interval $[\ell\kappa, (\ell + 1)\kappa)$, we assume the existence of a numerical method that simulates the estimated backward process $(\hat{y}_t)_{t \in [\ell\kappa, (\ell+1)\kappa]}$ with a bounded one-step numerical error.

Algorithm 1: Training Procedure of Denoising Markov Model

Input: Empirical distribution \widehat{p}_0 of the data distribution p_0 , the conditional probability $p_{t|0}(\cdot|x_0)$, number of total epochs N , batch size B , time distribution Ψ on $[0, T]$.

Output: The estimated backward generator \mathcal{K}_t^θ depending on model parameters θ through φ_t represented by a neural network.

1 **for** $n = 0$ **to** $N - 1$ **do**

2 Sample samples $\{x_0^{(k)}\}_{k=1}^B \sim \widehat{p}_0$ and time points $\{t^{(k)}\}_{k=1}^B \sim \Psi$ **for** $k = 1$ **to** B **do**

3 Sample $x_t^{(k)} \sim p_{t^{(k)}|0}(\cdot|x_0^{(k)})$ by simulating the forward process $(x_t)_{t \in [0, T]}$ with the generator $\mathcal{L}_{t^{(k)}}$

4 **end**

5 Compute the empirical average of the integrand in Eq. (3.5) as

$$\widehat{\mathfrak{L}}_{\text{SM}, \Psi}[\eta_{t|0}^\theta] = \frac{1}{B} \sum_{k=1}^B \left(\eta_{t^{(k)}|0}^\theta(\cdot|x_0^{(k)}) \mathcal{L}_{t^{(k)}} \eta_{t^{(k)}|0}^\theta(\cdot|x_0^{(k)})^{-1} + \mathcal{L}_{t^{(k)}} \log \eta_{t^{(k)}|0}^\theta(\cdot|x_0^{(k)}) \right) (x_t^{(k)});$$

Update the parameters θ by gradient descent as $\theta \leftarrow \theta - \epsilon \nabla \widehat{\mathfrak{L}}_{\text{SM}, \Psi}[\eta_{t|0}^\theta]$

6 **end**

Algorithm 2: Inference Procedure of Denoising Markov Model

Input: The estimated backward generator \mathcal{K}_t^θ depending on model parameters θ through φ_t represented by a neural network, a simple, easy-to-sample distribution $q_0 \approx p_T$.

Output: A sample \widehat{y}_T from $\widehat{q}_T \approx p_0$.

1 Sample $\widehat{y}_0 \sim \widehat{q}_0 = q_0$ **for** $\ell = 0$ **to** $L - 1$ **do**

2 Sample $\widehat{y}_{(\ell+1)\kappa}$ from the distribution $\widehat{q}_{(\ell+1)\kappa|\ell\kappa}(\cdot|\widehat{y}_{\ell\kappa})$ by the numerical method satisfying Assumption 3.10

3 **end**

Assumption 3.10 (One-Step Numerical Error) For each $\ell \in [0, L - 1]$, there exists an algorithm for simulating the estimated backward generator \mathcal{K}_t^θ . Given $\widehat{y}_{\ell\kappa}$, this algorithm induces a conditional distribution $\widehat{q}_{(\ell+1)\kappa|\ell\kappa}(\cdot|\widehat{y}_{\ell\kappa})$ approximating the estimated backward process $(y_t)_{t \in [\ell\kappa, (\ell+1)\kappa]}$ and the corresponding conditional distribution $q_{(\ell+1)\kappa|\ell\kappa}(\cdot|\widehat{y}_{\ell\kappa})$.

Moreover, we assume that the numerical error of the algorithm within each interval $[\ell\kappa, (\ell+1)\kappa]$ is of order $\mathcal{O}(\kappa^{1+r})$, i.e., the following bound holds:

$$\mathbb{E}_{\widehat{x}_{\ell\kappa} \sim \widehat{p}_{\ell\kappa}, \widehat{x}_{(\ell+1)\kappa} \sim \widehat{p}_{(\ell+1)\kappa}} \left[\log \frac{q_{(\ell+1)\kappa|\ell\kappa}(\widehat{x}_{(\ell+1)\kappa}|\widehat{x}_{\ell\kappa})}{\widehat{q}_{(\ell+1)\kappa|\ell\kappa}(\widehat{x}_{(\ell+1)\kappa}|\widehat{x}_{\ell\kappa})} \right] \lesssim \kappa^{1+r},$$

for some $r > 0$, where we write $a \lesssim b$ to mean $a \leq Cb$ for a constant $C > 0$ that may depend on fixed, time-uniform structural parameters of the model class (dimension, conditioning,

regularity), but is independent of the step size κ , the time horizon T , and any realized path.

Under this assumption, we obtain the following meta-error bound:

Theorem 3.11 (Meta-Error Bound for Denoising Markov Models) *Under Assumption 3.10, the distribution \hat{q}_T of the samples generated by the inference algorithm (Algorithm 2) satisfies:*

$$D_{\text{KL}}(p_0 \|\hat{q}_T) \lesssim D_{\text{KL}}(p_T \|\mathbf{q}_0) + \mathfrak{L}[\eta_t^\theta] + T\kappa^r, \quad (3.6)$$

where $\mathfrak{L}[\eta_t^\theta]$ is defined in Theorem 3.6.

The error bound (3.6) reveals three main sources of error in the model:

- (1) **Truncation Error** $D_{\text{KL}}(p_T \|\mathbf{q}_0)$: Results from approximating the forward process distribution p_T at the terminal time T by a simpler distribution \mathbf{q}_0 . If the forward process converges exponentially, as in certain cases of the continuous and discrete diffusion models discussed in Sections 2.2 and 2.3, this term typically becomes negligible for sufficiently large T .
- (2) **Estimation Error** $\mathfrak{L}[\eta_t^\theta]$: Corresponds exactly to the KL divergence between the path measure of the true backward process and that of its estimate as introduced in Theorem 3.6. This quantity measures how well φ_t aligns with the true marginal distributions p_t and can be directly minimized during training (Algorithm 1).
- (3) **Numerical Error** $T\kappa^r$: Represents cumulative numerical errors arising from the discrete approximation of the continuous-time backward process. This term heavily depends on the choice of numerical method and the underlying Markov process. Standard Euler-Maruyama methods typically yield $r = 1$. For general Lévy processes introduced in Section 4.3, the constant factor before the scaling $T\kappa^r$ depends on the dimension d , the conditioning, uniform regularity, and moment bounds on the coefficients of the Lévy triplet $(\vec{\mathbf{b}}_t^\varphi, \vec{\mathbf{D}}_t^\varphi, \vec{\lambda}_t^\varphi)$. Practitioners may choose adaptive discretization strategies based on the data distribution’s characteristics to reduce this error. For relevant discussions, we refer readers to Chen et al. (2023a); Benton et al. (2024a) (continuous diffusion models) and Chen and Ying (2025); Ren et al. (2025a) (discrete diffusion models).

4. Examples

In this section, we illustrate our framework by applying the theoretical results developed earlier to several important classes of Markov processes. Specifically, we revisit diffusion processes and jump processes, which recover continuous and discrete diffusion models as special cases. We then extend our analysis to general Lévy-type processes, revealing the broad applicability of our denoising Markov model framework. For clarity and brevity, we provide main results and intuitive explanations here, deferring rigorous mathematical details and proofs to Appendix B.

4.1 Diffusion Process

We first consider diffusion processes, where the underlying state space is the Euclidean space \mathbb{R}^d , and the generator acts on twice-differentiable functions, *i.e.*, functions in $C^2(\mathbb{R}^d)$.

Forward Process. We assume that the forward process $(\mathbf{x}_t)_{t \in [0, T]}$ is a diffusion process governed by the following SDE:

$$d\mathbf{x}_t = \mathbf{b}_t(\mathbf{x}_t)dt + \boldsymbol{\Sigma}_t(\mathbf{x}_t)d\mathbf{w}_t, \quad (4.1)$$

where $\mathbf{b}_t(\mathbf{x})$ is the drift, $(\mathbf{w}_t)_{t \geq 0}$ is the Wiener process, and the diffusion coefficient $\boldsymbol{\Sigma}_t(\mathbf{x})$ satisfies $\boldsymbol{\Sigma}_t(\mathbf{x})\boldsymbol{\Sigma}_t^\top(\mathbf{x}) = \mathbf{D}_t(\mathbf{x})$. Regularity conditions on the coefficients $\mathbf{b}_t(\mathbf{x})$ and $\mathbf{D}_t(\mathbf{x})$ are detailed in Assumption B.2.

The generator of the diffusion process $(\mathbf{x}_t)_{t \in [0, T]}$ is given by

$$\mathcal{L}_t f = \mathbf{b}_t(\mathbf{x}) \cdot \nabla f + \frac{1}{2} \mathbf{D}_t(\mathbf{x}) : \nabla^2 f,$$

where we adopt the notations $\nabla^2 = \nabla \nabla^\top$ and $\mathbf{D}_t(\mathbf{x}) : \nabla^2 f = \sum_{i,j=1}^d D_t^{ij}(\mathbf{x}) \partial_{ij} f$.

Backward Process. Given sufficient regularity (Assumption B.2), Theorem 3.2 shows that the backward generator $\tilde{\mathcal{L}}_t$ can be explicitly obtained as:

$$\tilde{\mathcal{L}}_{T-t} f = (-\mathbf{b}_t + \mathbf{D}_t \nabla \log p_t + \nabla \cdot \mathbf{D}_t) \cdot \nabla f + \frac{1}{2} \mathbf{D}_t : \nabla^2 f,$$

while our assumption (Assumption 3.3) leads to a similar structure for the estimated backward generator:

$$\mathcal{K}_{T-t} f = (-\mathbf{b}_t + \mathbf{D}_t \nabla \log \varphi_t + \nabla \cdot \mathbf{D}_t) \cdot \nabla f + \frac{1}{2} \mathbf{D}_t : \nabla^2 f,$$

implying that the estimated backward process $(\mathbf{y}_t)_{t \in [0, T]}$ is itself a diffusion process, described by the following SDE:

$$d\mathbf{y}_t = (-\mathbf{b}_{T-t}(\mathbf{y}_t) + \mathbf{D}_{T-t}(\mathbf{y}_t) \hat{\mathbf{s}}_{T-t}(\mathbf{y}_t) + \nabla \cdot \mathbf{D}_{T-t}(\mathbf{y}_t)) dt + \boldsymbol{\Sigma}_{T-t}(\mathbf{y}_t) d\mathbf{w}_t, \quad (4.2)$$

where $\hat{\mathbf{s}}_t^\theta = \nabla \log \varphi_t$ estimates the true score $\mathbf{s}_t = \nabla \log p_t$.

Loss Function. Applying Theorem 3.6, the KL divergence between the true and estimated backward processes is given by:

$$\mathfrak{L}[\hat{\mathbf{s}}_t] = \mathbb{E} \left[\int_0^T \frac{1}{2} \mathbf{D}_t(\mathbf{x}_t) : (\hat{\mathbf{s}}_t(\mathbf{x}_t) - \nabla \log p_t(\mathbf{x}_t)) (\hat{\mathbf{s}}_t(\mathbf{x}_t) - \nabla \log p_t(\mathbf{x}_t))^\top dt \right], \quad (4.3)$$

and the corresponding score-matching loss replaces $\nabla \log p_t$ by the conditional score $\nabla \log p_{t|0}$.

Example 4.1 (Continuous Diffusion Model) *Choosing the diffusion coefficient $\mathbf{D}_t(\mathbf{x}) = \mathbf{I}_d$ simplifies the score-matching loss to the familiar form from continuous diffusion models as in Eq. (2.2):*

$$\mathfrak{L}_{\text{SM}}[\hat{\mathbf{s}}_t] = \mathbb{E}_{\mathbf{x}_0 \sim p_0} \left[\int_0^T \mathbb{E}_{\mathbf{x}_t \sim p_{t|0}(\cdot|\mathbf{x}_0)} \left[\frac{1}{2} \|\hat{\mathbf{s}}_t(\mathbf{x}_t) - \nabla \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)\|^2 \right] dt \right],$$

matching exactly Eq. (2.2) up to time-reweighting.

4.2 Jump Process with Finite State Spaces

We now discuss finite-state jump processes, represented by continuous-time Markov chains on a finite state space \mathbb{X} , whose size we denote by $|\mathbb{X}|$.

Forward Process. The generator of the jump process on \mathbb{X} is given by

$$\mathcal{L}_t f(x) = \sum_{y \in \mathbb{X}} (f(y) - f(x)) \lambda_t(y, x),$$

where $\lambda_t(y, x)$ represents jump rates from x to y at time t , with $\lambda_t(x, x) = 0$ for all $x \in \mathbb{X}$. Then it is a classical result that the forward Markov chain $(x_t)_{t \in [0, T]}$ evolves according to (2.3) and one may verify that the jump rate $\lambda_t(y, x)$ matches that defined with the rate matrix $\mathbf{\Lambda}_t$ in Section 2.3.

Backward Process. Assumption 3.3 reduces to the following form:

$$\mathcal{K}_{T-t} f = \sum_{y \in \mathbb{X}} (f(y) - f(x)) \frac{\varphi_t(y)}{\varphi_t(x)} \lambda_t(x, y) = \sum_{y \in \mathbb{X}} (f(y) - f(x)) \widehat{s}_t(x, y) \lambda_t(x, y),$$

with estimated score function $\widehat{s}_t(x, y) = \varphi_t(y)/\varphi_t(x)$ approximating the true score $s_t(x, y) = p_t(y)/p_t(x)$. Thus, the estimated backward process $(y_t)_{t \in [0, T]}$ is also a continuous-time Markov chain with the form (2.4).

Loss Function. Theorem 3.6 now corresponds to

$$\mathfrak{L}[\widehat{s}_t] = \mathbb{E} \left[\int_0^T \sum_{y \in \mathbb{X}} \left(\frac{\widehat{s}_t(x_t, y)}{s_t(x_t, y)} - 1 - \log \frac{\widehat{s}_t(x_t, y)}{s_t(x_t, y)} \right) s_t(x_t, y) \lambda_t(x_t, y) dt \right]. \quad (4.4)$$

We adopt the convention $\lambda_t(x, x) = 0$ and therefore the summation is effectively excluding the case where $y = x_t$. We will also adopt this notation for brevity in the following discussions.

Example 4.2 (Discrete Diffusion Model) *As shown in Appendix B.2, the score-matching objective (3.5) in this case matches exactly the loss of discrete diffusion models (Eq. (2.5)) up to time reweighting.*

4.3 General Lévy-Type Process

We now investigate the most general setting for denoising Markov models, focusing on Lévy-type processes defined on \mathbb{R}^d . Lévy-type processes generalize both diffusion and jump processes and include a rich class of stochastic dynamics.

Forward Process. Under suitable assumptions on the generator (Assumption B.3), which ensure that the generator is sufficiently regular, *i.e.*, the domain of the augmented generator contains all compactly supported smooth functions, we characterize the forward generator explicitly using the Lévy-Khintchine representation:

Theorem 4.3 (General Form of Forward Process, Informal Version of Theorem B.4) *Under Assumption B.3, the generator \mathcal{L}_t of the forward process $(x_t)_{t \in [0, T]}$ can be represented in the following generalized Lévy-Khintchine form:*

$$\begin{aligned} \mathcal{L}_t f(\mathbf{x}) = & \mathbf{b}_t(\mathbf{x}) \cdot \nabla f(\mathbf{x}) + \frac{1}{2} \mathbf{D}_t(\mathbf{x}) : \nabla^2 f(\mathbf{x}) \\ & + \int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (f(\mathbf{y}) - f(\mathbf{x}) - (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) \mathbf{1}_{B(\mathbf{x}, 1)}(\mathbf{y} - \mathbf{x})) \lambda_t(d\mathbf{y}, \mathbf{x}), \end{aligned} \quad (4.5)$$

where $\mathbf{b}_t(\mathbf{x})$ is the drift term, $\mathbf{D}_t(\mathbf{x})$ is the diffusion coefficient, which is positive semidefinite for any $\mathbf{x} \in \mathbb{R}^d$, and $\lambda_t(d\mathbf{y}, \mathbf{x})$ is a Lévy measure, i.e., a Borel measure on $\mathbb{R}^d \setminus \{\mathbf{x}\}$ satisfying

$$\int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (1 \wedge \|\mathbf{y} - \mathbf{x}\|^2) \lambda_t(d\mathbf{y}, \mathbf{x}) < +\infty,$$

for any $\mathbf{x} \in \mathbb{R}^d$. The triple $(\mathbf{b}_t, \mathbf{D}_t, \lambda_t)$ is known as the Lévy triplet.

Note that if the coefficients $(\mathbf{b}_t, \mathbf{D}_t, \lambda_t)$ are time-invariant, this reduces to the classical Lévy-Khintchine form. By similar arguments as for (Çinlar and Jacod, 1981, Theorem 3.33), the generator \mathcal{L}_t of the general Lévy-type process $(\mathbf{x}_t)_{t \in \mathbb{R}}$ admits the stochastic integral representation (Proposition B.7):

$$\begin{aligned} \mathbf{x}_t = & \mathbf{x}_0 + \int_0^t \mathbf{b}_s(\mathbf{x}_s) ds + \int_0^t \boldsymbol{\Sigma}_s(\mathbf{x}_s) d\mathbf{w}_s \\ & + \int_0^t \int_{\mathbb{R}^d \setminus B(\mathbf{x}_{s-}, 1)} (\mathbf{y} - \mathbf{x}_{s-}) N[\lambda](ds, d\mathbf{y}) + \int_0^t \int_{B(\mathbf{x}_{s-}, 1)} (\mathbf{y} - \mathbf{x}_{s-}) \tilde{N}[\lambda](ds, d\mathbf{y}), \end{aligned}$$

where the matrix $\boldsymbol{\Sigma}_s(x)$ satisfies $\boldsymbol{\Sigma}_s(x) \boldsymbol{\Sigma}_s(x)^\top = \mathbf{D}_s(x)$ for any $x \in \mathbb{R}^d$, $(\mathbf{w}_s)_{s \geq 0}$ is a d -dimensional Wiener process, $N[\lambda](ds, d\mathbf{y})$ denotes the Poisson random measure with intensity $\lambda_s(d\mathbf{y}, \mathbf{x}_{s-})$, and $\tilde{N}[\lambda](ds, d\mathbf{y})$ is the compensated version of $N[\lambda](ds, d\mathbf{y})$, i.e.,

$$\tilde{N}[\lambda](ds, d\mathbf{y}) = N[\lambda](ds, d\mathbf{y}) - \lambda_s(d\mathbf{y}, \mathbf{x}_{s-}) ds.$$

Backward Process. Under further regularity conditions (Assumption B.6), particularly assuming that the Lévy measure $\lambda_t(d\mathbf{y}, \mathbf{x})$ admits a density, i.e., $\lambda_t(d\mathbf{y}, \mathbf{x}) = \lambda_t(\mathbf{y}, \mathbf{x}) d\mathbf{y}$, one may explicitly compute the backward generator $\tilde{\mathcal{L}}_t$, yielding the same form as Eq. (4.5):

$$\begin{aligned} \tilde{\mathcal{L}}_t f(\mathbf{x}) = & \tilde{\mathbf{b}}_t(\mathbf{x}) \cdot \nabla f(\mathbf{x}) + \frac{1}{2} \tilde{\mathbf{D}}_t(\mathbf{x}) : \nabla^2 f(\mathbf{x}) \\ & + \int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (f(\mathbf{y}) - f(\mathbf{x}) - (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) \mathbf{1}_{B(\mathbf{x}, 1)}(\mathbf{y} - \mathbf{x})) \tilde{\lambda}_t(d\mathbf{y}, \mathbf{x}), \end{aligned}$$

with modified Lévy triplet $(\tilde{\mathbf{b}}_t, \tilde{\mathbf{D}}_t, \tilde{\lambda}_t)$ explicitly given by:

$$\begin{aligned} \tilde{\mathbf{b}}_{T-t}(\mathbf{x}) = & -\mathbf{b}_t(\mathbf{x}) + \mathbf{D}_t(\mathbf{x}) \nabla \log p_t(\mathbf{x}) + \nabla \cdot \mathbf{D}_t(\mathbf{x}) + \mathcal{I}[p_t](\mathbf{x}), \\ \tilde{\mathbf{D}}_{T-t}(\mathbf{x}) = & \mathbf{D}_t(\mathbf{x}), \quad \text{and} \quad \tilde{\lambda}_{T-t}(\mathbf{y}, \mathbf{x}) = \frac{p_t(\mathbf{y})}{p_t(\mathbf{x})} \lambda_t(\mathbf{x}, \mathbf{y}), \end{aligned}$$

where the integral term $\mathcal{I}[p_t]$ accounts for asymmetric jumps:

$$\mathcal{I}[p_t](\mathbf{x}) = \int_{B(\mathbf{x},1)} (\mathbf{y} - \mathbf{x}) \left(\frac{p_t(\mathbf{y})}{p_t(\mathbf{x})} \lambda_t(\mathbf{x}, \mathbf{y}) + \lambda_t(\mathbf{y}, \mathbf{x}) \right) d\mathbf{y}.$$

Under our estimation assumption (Assumption 3.3), the estimated backward generator \mathcal{K}_t shares the same general Lévy-type form (4.5) but with the estimated Lévy triplet $(\tilde{\mathbf{b}}_t^\varphi, \tilde{\mathbf{D}}_t^\varphi, \tilde{\lambda}_t^\varphi)$:

$$\begin{aligned} \tilde{\mathbf{b}}_{T-t}^\varphi(\mathbf{x}) &= -\mathbf{b}_t(\mathbf{x}) + \mathbf{D}_t(\mathbf{x}) \nabla \log \varphi_t(\mathbf{x}) + \nabla \cdot \mathbf{D}_t(\mathbf{x}) + \mathcal{I}[\varphi_t](\mathbf{x}) \\ \tilde{\mathbf{D}}_{T-t}^\varphi(\mathbf{x}) &= \mathbf{D}_t(\mathbf{x}), \quad \text{and} \quad \tilde{\lambda}_{T-t}^\varphi(\mathbf{y}, \mathbf{x}) = \frac{\varphi_t(\mathbf{y})}{\varphi_t(\mathbf{x})} \lambda_t(\mathbf{x}, \mathbf{y}). \end{aligned}$$

We remark that the integral $\mathcal{I}[\cdot]$ should be understood as a *bona fide* integral or in the Cauchy principal value sense, depending on the singularity of $\lambda_t(\mathbf{y}, \mathbf{x})$ (see Conforti and Léonard (2022, Theorem 5.7) for sufficient conditions and Example 4.4 for an explicit example).

Loss Function. To practically implement this model, we define the diffusion and jump score functions separately as:

$$\mathbf{s}_t^{\text{diff}}(\mathbf{x}) = \nabla \log p_t(\mathbf{x}), \quad \text{and} \quad \hat{\mathbf{s}}_t^{\text{diff}}(\mathbf{x}) = \nabla \log \varphi_t(\mathbf{x}),$$

and

$$s_t^{\text{jump}}(\mathbf{x}, \mathbf{y}) = \frac{p_t(\mathbf{y})}{p_t(\mathbf{x})}, \quad \text{and} \quad \hat{s}_t^{\text{jump}}(\mathbf{x}, \mathbf{y}) = \frac{\varphi_t(\mathbf{y})}{\varphi_t(\mathbf{x})},$$

respectively. Using these definitions, the KL divergence between the backward processes decomposes naturally into diffusion and jump parts, yielding the unified loss function:

$$\begin{aligned} \mathfrak{L}[\hat{\mathbf{s}}_t^{\text{diff}}, \hat{s}_t^{\text{jump}}] &= \mathbb{E} \left[\int_0^T \frac{1}{2} \mathbf{D}_t(\mathbf{x}_t) : \left(\hat{\mathbf{s}}_t^{\text{diff}}(\mathbf{x}_t) - \nabla \log p_t(\mathbf{x}_t) \right) \left(\hat{\mathbf{s}}_t^{\text{diff}}(\mathbf{x}_t) - \nabla \log p_t(\mathbf{x}_t) \right)^\top dt \right. \\ &\quad \left. + \int_0^T \int_{\mathbb{R}^d} \left(\frac{\hat{s}_t^{\text{jump}}(\mathbf{x}_t, \mathbf{y})}{s_t^{\text{jump}}(\mathbf{x}_t, \mathbf{y})} - 1 - \log \frac{\hat{s}_t^{\text{jump}}(\mathbf{x}_t, \mathbf{y})}{s_t^{\text{jump}}(\mathbf{x}_t, \mathbf{y})} \right) s_t^{\text{jump}}(\mathbf{x}_t, \mathbf{y}) \lambda_t(\mathbf{x}_t, \mathbf{y}) d\mathbf{y} dt \right], \end{aligned} \quad (4.6)$$

which is exactly the sum of the KL divergence (4.3) derived from the discussion of diffusion processes and that of the jump process (4.4), and the corresponding score-matching loss is thus

$$\begin{aligned} \mathfrak{L}_{\text{SM}}[\hat{\mathbf{s}}_t^{\text{diff}}, \hat{s}_t^{\text{jump}}] &= \mathbb{E}_{\mathbf{x}_0 \sim p_0} \left[\int_0^T \mathbb{E}_{\mathbf{x}_t \sim p_{t|0}(\cdot|\mathbf{x}_0)} \left[\frac{1}{2} \mathbf{D}_t(\mathbf{x}_t) : \left(\hat{\mathbf{s}}_t^{\text{diff}}(\mathbf{x}_t) - \nabla \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0) \right) \left(\hat{\mathbf{s}}_t^{\text{diff}}(\mathbf{x}_t) - \nabla \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0) \right)^\top \right. \right. \\ &\quad \left. \left. + \int_{\mathbb{R}^d} \left(\hat{s}_t^{\text{jump}}(\mathbf{x}_t, \mathbf{y}) - \frac{p_{t|0}(\mathbf{y}|\mathbf{x}_0)}{p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)} \log \hat{s}_t^{\text{jump}}(\mathbf{x}_t, \mathbf{y}) \right) \lambda_t(\mathbf{x}_t, \mathbf{y}) d\mathbf{y} \right] dt \right]. \end{aligned} \quad (4.7)$$

Example 4.4 (Lévy-Itô Model (Yoon et al., 2023)) Consider the special case where the diffusion component vanishes, i.e., $\mathbf{D}_t(\mathbf{x}) \equiv 0$, and the Lévy measure $\lambda_t(\mathbf{y}, \mathbf{x})d\mathbf{y}$ is chosen as a time-homogeneous and isotropic α -stable Lévy measure with the following form:

$$\lambda_t(\mathbf{y}, \mathbf{x})d\mathbf{y} = \frac{\alpha 2^{\alpha-1} \pi^{-d/2} \Gamma(\frac{\alpha+d}{2})}{\Gamma(1 - \frac{\alpha}{2})} \frac{1}{\|\mathbf{y} - \mathbf{x}\|^{d+\alpha}} d\mathbf{y}. \quad (4.8)$$

In this setting, the integral term $\mathcal{I}[\cdot]$ arising in the backward drift admits a concise expression in the Fourier domain (Yoon et al., 2023, Theorem B.3):

$$\mathcal{I}[f](\mathbf{x}) = \frac{1}{f(\mathbf{x})} \mathcal{F}^{-1} \left\{ -i\alpha \boldsymbol{\xi} \|\boldsymbol{\xi}\|^{\alpha-2} \mathcal{F}\{f\} \right\}(\mathbf{x}),$$

where \mathcal{F} and \mathcal{F}^{-1} denote the Fourier transform and its inverse, respectively.

While the Lévy measures associated with the true and estimated backward processes differ from the forward measure, a practical simplification used by Yoon et al. (2023) is to approximate the ratio terms $p_t(\mathbf{y})/p_t(\mathbf{x})$ or $\varphi_t(\mathbf{y})/\varphi_t(\mathbf{x})$ by 1. This approximation is justified by noting that the Lévy measure defined in Eq. (4.8) concentrates strongly near the origin, thus allowing the standard isotropic α -stable process to be used directly in the backward simulation.

Typically, to obtain the estimated drift term $\tilde{\mathbf{b}}_t^\varphi(\mathbf{x})$, one would first estimate the jump-related score function $\hat{s}_t^{\text{jump}}(\mathbf{x}, \mathbf{y})$ by optimizing the general Lévy-type score-matching loss (4.7), and then plug this into the integral term $\mathcal{I}[\varphi_t](\mathbf{x})$. However, an alternative approach proposed by Yoon et al. (2023) is to directly learn the integral term $\mathcal{I}[\varphi_t](\mathbf{x})$ by minimizing the following simpler loss:

$$\mathfrak{L}_{\text{LIM}}[\mathcal{I}[\varphi_t]] = \mathbb{E}_{\mathbf{x}_0 \sim p_0} \left[\int_0^T \mathbb{E}_{\mathbf{x}_t \sim p_{t|0}(\cdot|\mathbf{x}_0)} \left[\|\mathcal{I}[\varphi_t](\mathbf{x}) - \mathcal{I}[p_{t|0}(\cdot|\mathbf{x}_0)](\mathbf{x})\|^2 \right] dt \right].$$

This loss no longer directly corresponds to minimizing an upper bound on $D_{\text{KL}}(p_0 \| q_T)$ as in Corollary 3.7, but it is practically appealing due to its explicit form. When the forward drift is chosen as $\mathbf{b}_t(\mathbf{x}) = -\mathbf{x}/\alpha$, this integral term simplifies to a closed-form expression:

$$\mathcal{I}[p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)](\mathbf{x}) = -\frac{\mathbf{x}_t - e^{-\frac{t}{\alpha}} \mathbf{x}_0}{\alpha(1 - e^{-t})},$$

thus enabling efficient training and inference.

In the following theorem we justify the necessity of allowing general Lévy-type forward processes in the denoising Markov model framework, including instances such as the Lévy-Itô model (Yoon et al., 2023). We show that, under the standard diffusion-only setting with affine drift and additive Gaussian noise, heavy-tailed data inevitably induce an *infinite* terminal mismatch against a Gaussian reference. Consequently, the KL-based error bound degenerates, which motivates enlarging the design space beyond diffusions.

Theorem 4.5 Consider a denoising Markov model whose forward process is the diffusion in Eq. (4.1),

$$d\mathbf{x}_t = \mathbf{b}_t(\mathbf{x}_t)dt + \boldsymbol{\Sigma}_t d\mathbf{w}_t,$$

with affine drift $\mathbf{b}_t(\mathbf{x}) = \mathbf{A}_t\mathbf{x} + \mathbf{u}_t$, where $\mathbf{A}_t \in \mathbb{R}^{d \times d}$ and $\mathbf{u}_t \in \mathbb{R}^d$ are measurable and bounded on $[0, T]$, and Σ_t is measurable and bounded on $[0, T]$. Assume the data distribution p_0 is absolutely continuous, spherically symmetric, and heavy-tailed with $\mathbb{E}[\|\mathbf{x}_0\|^2] = \infty$, and let the reference q_0 be Gaussian. Then, for every time horizon $T > 0$,

$$D_{\text{KL}}(p_T \| q_0) = \infty,$$

where p_T is the T -marginal of the forward process.

The proof is provided in Appendix B.3. By Corollary 3.7, the sample error bound degenerates for diffusion-only forward processes with affine drift and additive Gaussian noise when p_0 is heavy-tailed:

$$D_{\text{KL}}(p_0 \| q_T) \leq D_{\text{KL}}(p_T \| q_0) + \mathfrak{L}[\eta_t] = \infty,$$

because the terminal marginal p_T retains heavy tails, making the terminal mismatch against a Gaussian reference q_0 infinite and the KL-based guarantee vacuous.

Remark 4.6 *The failure mode of Theorem 4.5 is intrinsic to any finite physical time Gaussian noising path and therefore cannot be eliminated by a mere time reparametrization. Consider compressing $t \in [0, \infty)$ into $\tau \in [0, 1)$ via an increasing C^1 map $g : [0, 1) \rightarrow [0, \infty)$ with $g(\tau) \rightarrow \infty$ as $\tau \rightarrow 1$, and the reparametrized process $\mathbf{y}_\tau := \mathbf{x}_{g(\tau)}$. The time-changed process satisfies*

$$d\mathbf{y}_\tau = g'(\tau)\mathbf{b}_{g(\tau)}(\mathbf{y}_\tau)d\tau + \sqrt{g'(\tau)}\Sigma_{g(\tau)}d\mathbf{w}_\tau,$$

and since $g(\tau) = \int_0^\tau g'(s)ds \rightarrow \infty$ as $\tau \rightarrow 1$, the factor g' cannot remain bounded and the dynamics become stiff near $\tau \rightarrow 1$. As a result, any implementable numerical scheme must effectively truncate at $\tau = 1 - \varepsilon$ for some $\varepsilon > 0$ (or equivalently, at finite physical time $t = g(1 - \varepsilon)$), where the KL divergence remains infinite.

This exposes a concrete failure mode of diffusion-only denoising Markov models for heavy-tailed data and motivates enlarging the design space to Lévy-type forward processes with heavy-tailed marginals, *e.g.*, the Ornstein-Uhlenbeck process driven by an α -stable Lévy process, whose stationary law is α -stable. In such cases one can simply choose this heavy-tailed stationary distribution as the reference q_0 to match the tails, restoring a finite and informative terminal bound and making Corollary 3.7 meaningful. For empirical evidence on using α -stable Lévy forwards with real-world heavy-tailed data, we refer to the results in (Yoon et al., 2023), without repeating those experiments in this paper. This provides a principled rationale for adopting general Lévy-type processes in our framework when heavy-tailed data may arise. This framework is practically attractive when $p_{t|0}$, or a surrogate such as $\mathcal{I}[p_{t|0}]$, is available in closed form or can be approximated efficiently.

To end this section, we provide insights into the connections between denoising Markov models and generator matching (Holderrieth et al., 2025) in the following remark.

Remark 4.7 (Connections to Generator Matching (Holderrieth et al., 2025))

The primary conceptual distinction between our proposed framework of denoising Markov models and the generator matching approach of Holderrieth et al. (2025) lies in the direction and strategy used to specify and estimate the underlying Markov dynamics.

In denoising Markov models, we specify the forward generator \mathcal{L}_t explicitly and subsequently estimate the backward generator $\tilde{\mathcal{L}}_t$ through the parametrization of a function φ_t , guided by conditional densities $p_{t|0}$. In contrast, the generator matching method specifies the form of the backward generator $\tilde{\mathcal{L}}_t$ first and then directly learns its coefficients using a prescribed probability path $(p_t)_{t \in [0, T]}$ (or equivalently, $(\tilde{p}_t)_{t \in [0, T]}$).

Specifically, generator matching assumes that for each terminal condition $\tilde{\mathbf{x}}_T \sim \tilde{p}_T$ (or equivalently, $\mathbf{x}_0 \sim p_0$), the conditional probability path $(\tilde{p}_{t|T}(\cdot|\tilde{\mathbf{x}}_T))_{t \in [0, T]}$ (or equivalently, $(p_{T-t|0}(\cdot|\mathbf{x}_0))_{t \in [0, T]}$) is pre-determined and assumed to be governed by a general Lévy-type conditional backward generator of the form:

$$\begin{aligned} \tilde{\mathcal{L}}_{t|T}^{\tilde{\mathbf{x}}_T} &= \tilde{\mathbf{b}}_{t|T}(\mathbf{x}|\tilde{\mathbf{x}}_T) \cdot \nabla f(\mathbf{x}) + \frac{1}{2} \tilde{\mathbf{D}}_{t|T}(\mathbf{x}|\tilde{\mathbf{x}}_T) : \nabla^2 f(\mathbf{x}) \\ &+ \int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (f(\mathbf{y}) - f(\mathbf{x}) - (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) \mathbf{1}_{B(\mathbf{x}, 1)}(\mathbf{y} - \mathbf{x})) \tilde{\lambda}_{t|T}(d\mathbf{y}, \mathbf{x}|\tilde{\mathbf{x}}_T), \end{aligned} \quad (4.9)$$

where the triplet $(\tilde{\mathbf{b}}_{t|T}(\mathbf{x}|\tilde{\mathbf{x}}_T), \tilde{\mathbf{D}}_{t|T}(\mathbf{x}|\tilde{\mathbf{x}}_T), \tilde{\lambda}_{t|T}(d\mathbf{y}, \mathbf{x}|\tilde{\mathbf{x}}_T))$ is assumed to be available in closed form formulas, derived directly from Kolmogorov's forward equation (Theorem A.23):

$$\partial_t \tilde{p}_{t|T}(\cdot|\tilde{\mathbf{x}}_T) = \left(\tilde{\mathcal{L}}_{t|T}^{\tilde{\mathbf{x}}_T} \right)^* \tilde{p}_{t|T}(\cdot|\tilde{\mathbf{x}}_T).$$

Then, leveraging the linearity of the generator, the unconditional backward generator $\tilde{\mathcal{L}}_t$ can be expressed as the posterior expectation of this conditional generator (cf. (Holderrrieth et al., 2025, Proposition 1)):

$$\begin{aligned} \tilde{\mathcal{L}}_t f(\mathbf{x}) &= \mathbb{E}_{\tilde{\mathbf{x}}_T \sim \tilde{p}_{T|t}(\cdot|\mathbf{x})} \left[\tilde{\mathcal{L}}_{t|T}^{\tilde{\mathbf{x}}_T} f(\mathbf{x}) \right] \\ &= \mathbb{E}_{\tilde{\mathbf{x}}_T \sim \tilde{p}_{T|t}(\cdot|\mathbf{x})} \left[\tilde{\mathbf{b}}_{t|T}(\mathbf{x}|\tilde{\mathbf{x}}_T) \right] \cdot \nabla f(\mathbf{x}) + \frac{1}{2} \mathbb{E}_{\tilde{\mathbf{x}}_T \sim \tilde{p}_{T|t}(\cdot|\mathbf{x})} \left[\tilde{\mathbf{D}}_{t|T}(\mathbf{x}|\tilde{\mathbf{x}}_T) \right] : \nabla^2 f(\mathbf{x}) \\ &+ \int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (f(\mathbf{y}) - f(\mathbf{x}) - (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) \mathbf{1}_{B(\mathbf{x}, 1)}(\mathbf{y} - \mathbf{x})) \mathbb{E}_{\tilde{\mathbf{x}}_T \sim \tilde{p}_{T|t}(\cdot|\mathbf{x})} \left[\tilde{\lambda}_{t|T}(d\mathbf{y}, \mathbf{x}|\tilde{\mathbf{x}}_T) \right], \end{aligned} \quad (4.10)$$

i.e., the Lévy triplet $(\tilde{\mathbf{b}}_t, \tilde{\mathbf{D}}_t, \tilde{\lambda}_t)$ of the backward process is given by the posterior expectations of the Lévy triplet $(\tilde{\mathbf{b}}_{t|T}(\cdot|\tilde{\mathbf{x}}_T), \tilde{\mathbf{D}}_{t|T}(\cdot|\tilde{\mathbf{x}}_T), \tilde{\lambda}_{t|T}(\cdot|\tilde{\mathbf{x}}_T))$ of the conditional generator $\tilde{\mathcal{L}}_{t|T}^{\tilde{\mathbf{x}}_T}$ given $\tilde{\mathbf{x}}_T$, as

$$\begin{aligned} \tilde{\mathbf{b}}_t(\mathbf{x}) &= \mathbb{E}_{\tilde{\mathbf{x}}_T \sim \tilde{p}_{T|t}(\cdot|\mathbf{x})} \left[\tilde{\mathbf{b}}_{t|T}(\mathbf{x}|\tilde{\mathbf{x}}_T) \right], \quad \tilde{\mathbf{D}}_t(\mathbf{x}) = \mathbb{E}_{\tilde{\mathbf{x}}_T \sim \tilde{p}_{T|t}(\cdot|\mathbf{x})} \left[\tilde{\mathbf{D}}_{t|T}(\mathbf{x}|\tilde{\mathbf{x}}_T) \right], \\ \text{and } \tilde{\lambda}_t(d\mathbf{y}, \mathbf{x}) &= \mathbb{E}_{\tilde{\mathbf{x}}_T \sim \tilde{p}_{T|t}(\cdot|\mathbf{x})} \left[\tilde{\lambda}_{t|T}(d\mathbf{y}, \mathbf{x}|\tilde{\mathbf{x}}_T) \right]. \end{aligned}$$

The generator matching training strategy thus involves learning these posterior expectations (4.10) by formulating suitable regression problems. For instance, the drift term $\tilde{\mathbf{b}}_t(\mathbf{x})$ can be learned via the minimization of the following regression-based loss:

$$\begin{aligned} \mathfrak{L}_{\text{GM}}^{\mathbf{b}}[\tilde{\mathbf{b}}_t^\theta] &= \int_0^T \mathbb{E}_{\tilde{\mathbf{x}}_t \sim \tilde{p}_t} \left[\mathbb{E}_{\tilde{\mathbf{x}}_T \sim \tilde{p}_{T|t}(\cdot|\tilde{\mathbf{x}}_t)} \left[\left\| \tilde{\mathbf{b}}_t^\theta(\tilde{\mathbf{x}}_t) - \tilde{\mathbf{b}}_{t|T}(\tilde{\mathbf{x}}_t|\tilde{\mathbf{x}}_T) \right\|^2 \right] \right] dt \\ &= \int_0^T \mathbb{E}_{\mathbf{x}_0 \sim p_0} \left[\mathbb{E}_{\mathbf{x}_{T-t} \sim p_{T-t|0}(\cdot|\mathbf{x}_0)} \left[\left\| \tilde{\mathbf{b}}_t^\theta(\mathbf{x}_{T-t}) - \tilde{\mathbf{b}}_{t|T}(\mathbf{x}_{T-t}|\mathbf{x}_0) \right\|^2 \right] \right] dt, \end{aligned}$$

with samples $\tilde{\mathbf{x}}_t \sim \tilde{p}_{t|T}(\cdot|\tilde{\mathbf{x}}_T)$ obtained by simulating the conditional generator (4.9). After obtaining the estimated Lévy triplet $(\tilde{\mathbf{b}}_t^\theta, \tilde{\mathbf{D}}_t^\theta, \tilde{\lambda}_t^\theta)$, inference proceeds via simulation from the resulting estimated backward generator.

5. Experiments

In this section, we present several experiments validating our theoretical results. Given the extensive empirical literature on continuous and discrete diffusion models as special cases of our framework (*cf.* Examples 4.1 and 4.2), we focus on novel examples of Markov processes that have not yet been extensively explored in the context of denoising Markov models. The primary purpose of these experiments is to demonstrate the generality and flexibility of our theoretical approach. Given the established success of diffusion-based models, our objective is not to outperform existing methods; instead, we aim to highlight the viability of alternative formulations by applying them to conceptually simple tasks.

We will consider two examples of denoising Markov models: (1) a diffusion model with geometric Brownian motion as the forward process, and (2) a “jump model” with a jump process in \mathbb{R}^d as the forward process. In both cases, we will demonstrate that the learned backward processes can effectively approximate the target distribution.

5.1 Diffusion Model with Geometric Brownian Motion

In this illustrative example, we consider geometric Brownian motion (GBM) as our forward process. The forward dynamics follow:

$$d\mathbf{x}_t = \mathbf{x}_t \odot \Sigma d\mathbf{w}_t,$$

where \odot denotes the element-wise product. Through a change of variables, this equation admits a closed-form solution:

$$\mathbf{x}_t = \mathbf{x}_0 \odot \exp\left(\Sigma \mathbf{w}_t - \frac{1}{2} \text{diag}(\Sigma \Sigma^\top) t\right), \quad (5.1)$$

where $\text{diag}(\Sigma \Sigma^\top)$ denotes the vector with the diagonal elements of the matrix $\Sigma \Sigma^\top$. The diffusion matrix in this case is given explicitly by $\mathbf{D}(\mathbf{x}_t) = \text{diag}(\mathbf{x}_t) \Sigma \Sigma^\top \text{diag}(\mathbf{x}_t)$. Thus, this choice clearly fits within our denoising Markov model framework, with a diffusion process as the forward dynamics. Unlike the standard diffusion models (Example 4.1), the diffusion matrix here is neither constant nor spatially homogeneous, depending explicitly on the current state \mathbf{x}_t .

Following the framework in Section 4.1, we parametrize the estimated score function $\hat{\mathbf{s}}_t(\mathbf{x})$ by a neural network with parameters θ and optimize the following score-matching loss (4.3):

$$\mathcal{L}_{\text{SM}}[\hat{\mathbf{s}}_t] = \mathbb{E}_{\mathbf{x}_0 \sim p_0} \left[\int_0^T \mathbb{E}_{\mathbf{x}_t \sim p_{t|0}(\cdot|\mathbf{x}_0)} \left[\frac{1}{2} \mathbf{D}(\mathbf{x}_t) : \left(\hat{\mathbf{s}}_t^\theta(\mathbf{x}_t) - \nabla \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0) \right) \left(\hat{\mathbf{s}}_t^\theta(\mathbf{x}_t) - \nabla \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0) \right)^\top \right] dt \right].$$

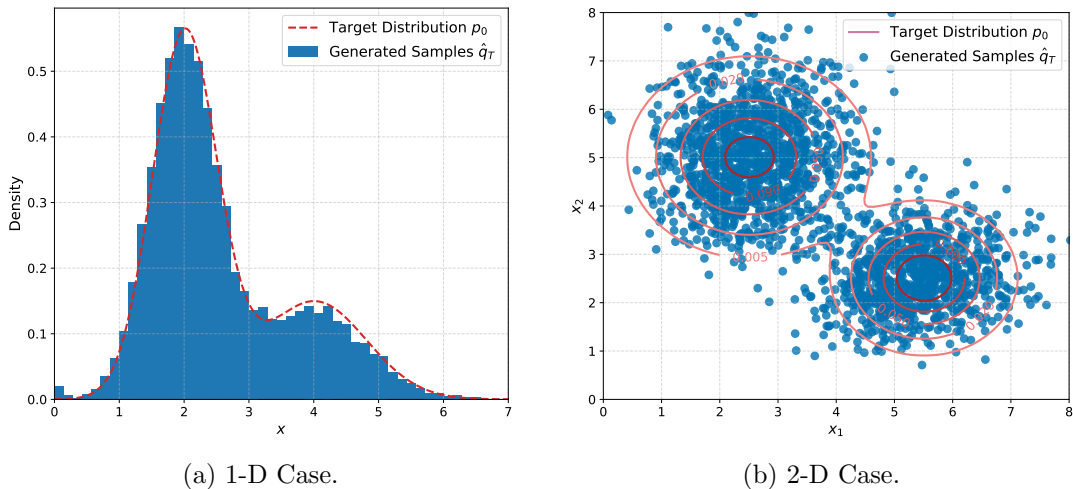


Figure 2: Empirical results of denoising Markov models using geometric Brownian motion as the forward process. The generated samples from \hat{q}_T (in blue) closely match the target distribution p_0 (dashed line in 1-D and contour lines in 2-D in red), highlighting the effectiveness of the learned backward process.

Using the explicit solution (5.1), the gradient of the conditional log-density is given by:

$$\nabla \log p_{t|0}(\mathbf{x}_t | \mathbf{x}_0) = -\mathbf{x}_t^{-1} \odot \left(1 + \frac{1}{t} (\boldsymbol{\Sigma} \boldsymbol{\Sigma}^\top)^{-1} \left(\log \mathbf{x}_t - \log \mathbf{x}_0 + \frac{t}{2} \text{diag} \left(\boldsymbol{\Sigma} \boldsymbol{\Sigma}^\top \right) \right) \right), \quad (5.2)$$

with the vector operations taken element-wise. Note that, when $\mathbf{D}(\mathbf{x}_t)$ is not diagonal, this loss no longer simplifies to a standard mean squared error, highlighting the generality of our method.

Due to the positivity-preserving nature of GBM, our experiments focus on the positive orthant \mathbb{R}_+^d . We choose the target distribution as the absolute value of mixtures of Gaussian distributions. We select a sufficiently large terminal time T to ensure that the forward dynamics (5.1) approach the origin closely with high probability. Correspondingly, the backward dynamics, derived from Eq. (4.2), read:

$$d\mathbf{y}_t = (\mathbf{D}(\mathbf{y}_t) \hat{\mathbf{s}}_{T-t}(\mathbf{y}_t) + \nabla \cdot \mathbf{D}(\mathbf{y}_t)) dt + \mathbf{y}_t \odot \boldsymbol{\Sigma} d\mathbf{w}_t.$$

To demonstrate the effectiveness of our method, we conduct experiments in both one and two-dimensional cases (Figure 2). For both cases, we use an MLP with 5 layers and 128 hidden units for the score function $\hat{\mathbf{s}}_t(\mathbf{x})$ with the input simply being the concatenation of the state \mathbf{x}_t and the time t . We optimize the score-matching loss (4.3) using the Adam optimizer with a learning rate of 10^{-3} . In the 1-D scenario, we set $\boldsymbol{\Sigma} = 1$ and use the following target distribution:

$$p_0(y) = 0.7\mathcal{N}(y; 2, 0.25) + 0.3\mathcal{N}(y; 4, 0.64), \quad \forall y \in \mathbb{R}_+.$$

In the 2D scenario, we use:

$$\boldsymbol{\Sigma} = \begin{bmatrix} 1.0 & 0.4 \\ 0.2 & 1.0 \end{bmatrix}, \quad \text{and } p_0(\mathbf{y}) = 0.6\mathcal{N}\left(\mathbf{y}; \begin{bmatrix} 2.5 \\ 5 \end{bmatrix}, \mathbf{I}\right) + 0.4\mathcal{N}\left(\mathbf{y}; \begin{bmatrix} 5.5 \\ 2.5 \end{bmatrix}, \mathbf{I}\right), \quad \forall \mathbf{y} \in \mathbb{R}_+^2,$$

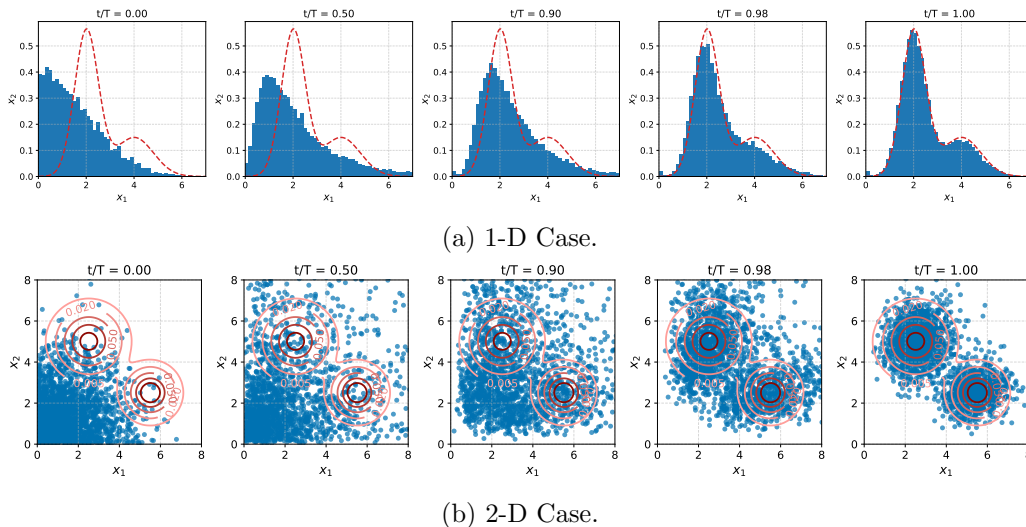


Figure 3: Visualization of generated backward trajectories using geometric Brownian motion as the forward process. Samples generated via the learned backward dynamics (in blue) evolve from the initial distribution q_0 with the estimated score function $\widehat{\mathbf{s}}_t(\mathbf{x})$ into the target distribution p_0 (in red), further validating the theoretical robustness and practical effectiveness of our denoising Markov models.

where $\mathcal{N}(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the Gaussian density with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. We approximate the initial distribution q_0 by sampling from $\mathcal{N}(\mathbf{0}, 2\mathbf{I})$ and then applying the absolute-value transformation.

As evidenced in Figure 2, our model achieves satisfactory performance: the generated samples closely approximate the target distributions in both dimensions. Despite the potential singularities in the score function (5.2) as $t \rightarrow 0$ or as $\|\mathbf{x}_t\| \rightarrow 0$, training is notably stable, requiring only moderate neural-network sizes. Furthermore, Figure 3 illustrates the evolution of the generated trajectories during inference, demonstrating how samples move smoothly from the initial noise toward the complex target distribution along a different path from the pure Brownian motion (Variance-Exploding SDE) or OU process (Variance-Preserving SDE). This provides empirical support for the soundness of the backward process learned via score-matching and emphasizes our framework’s flexibility in accommodating more general Markov dynamics beyond standard diffusion processes.

5.2 Jump Model

In this example, we propose a pure jump process-based model as a new type of denoising Markov model. Specifically, we consider a forward process consisting exclusively of jumps, *i.e.*, we set the drift $\mathbf{b}_t(\mathbf{x}) \equiv \mathbf{0}$ and the diffusion coefficient $\mathbf{D}_t(\mathbf{x}) \equiv \mathbf{0}$. For simplicity, we assume that the Lévy measure $\lambda_t(d\mathbf{y}, \mathbf{x})$ admits a density $\lambda_t(\mathbf{y}, \mathbf{x})$ with respect to the Lebesgue measure and satisfies the following integrability condition:

$$\int_{B(\mathbf{x}, 1)} \|\mathbf{y} - \mathbf{x}\| \lambda_t(d\mathbf{y}, \mathbf{x}) < +\infty.$$

This ensures that small jumps in Eq. (4.5) are integrable and thus can be naturally included with the drift term, which vanishes in this case.

Under these assumptions, the forward generator simplifies to:

$$\mathcal{L}_t f(\mathbf{x}) = \int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (f(\mathbf{y}) - f(\mathbf{x})) \lambda_t(\mathbf{y}, \mathbf{x}) d\mathbf{y}.$$

This corresponds to a pure-jump Lévy-type process represented by the stochastic integral:

$$\mathbf{x}_t = \mathbf{x}_0 + \int_0^t \int_{\mathbb{R}^d} (\mathbf{y} - \mathbf{x}_{s-}) N[\lambda](ds, d\mathbf{y}), \quad (5.3)$$

where $N[\lambda](ds, d\mathbf{y})$ denotes the Poisson random measure with compensator $\lambda_s(\mathbf{y}, \mathbf{x}_{s-}) d\mathbf{y} ds$. We refer readers to Section 4.2 for more details on this formulation.

Correspondingly, Assumption 3.3 yields a simplified estimated backward generator:

$$\mathcal{K}_t f(\mathbf{x}) = \int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (f(\mathbf{y}) - f(\mathbf{x})) \tilde{\lambda}_t^\varphi(\mathbf{y}, \mathbf{x}) d\mathbf{y},$$

where the estimated Lévy measure $\tilde{\lambda}_t^\varphi(\mathbf{y}, \mathbf{x}) d\mathbf{y}$ is defined explicitly via the score function for the jump part:

$$\tilde{\lambda}_t^\varphi(\mathbf{y}, \mathbf{x}) = \hat{s}_{T-t}^{\text{jump}}(\mathbf{x}, \mathbf{y}) \lambda_{T-t}(\mathbf{x}, \mathbf{y}) = \frac{\varphi_{T-t}(\mathbf{y})}{\varphi_{T-t}(\mathbf{x})} \lambda_{T-t}(\mathbf{x}, \mathbf{y}).$$

This estimated Lévy measure is both temporally and spatially inhomogeneous. In contrast to general Lévy-type processes, this jump-only model eliminates the need to estimate any diffusion-based score function. The score-matching loss from (4.7) then reduces to:

$$\begin{aligned} \mathfrak{L}_{\text{SM}}[\hat{s}_t^{\text{jump}}] = & \mathbb{E}_{\mathbf{x}_0 \sim p_0} \left[\int_0^T \mathbb{E}_{\mathbf{x}_t \sim p_{t|0}(\cdot|\mathbf{x}_0)} \left[\int_{\mathbb{R}^d} \left(\hat{s}_t^{\text{jump}}(\mathbf{x}_t, \mathbf{y}) - \frac{p_{t|0}(\mathbf{y}|\mathbf{x}_0)}{p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)} \log \hat{s}_t^{\text{jump}}(\mathbf{x}_t, \mathbf{y}) \right) \lambda_t(\mathbf{x}_t, \mathbf{y}) d\mathbf{y} \right] dt \right]. \end{aligned} \quad (5.4)$$

Once the jump-based score function \hat{s}_t^{jump} is learned, we simulate the backward process by:

$$\mathbf{y}_t = \mathbf{y}_0 + \int_0^t \int_{\mathbb{R}^d} (\mathbf{y} - \mathbf{y}_{s-}) N[\tilde{\lambda}^\varphi](ds, d\mathbf{y}), \quad (5.5)$$

where $N[\tilde{\lambda}^\varphi](ds, d\mathbf{y})$ denotes a Poisson random measure with compensator $\tilde{\lambda}_t^\varphi(\mathbf{y}, \mathbf{y}_{s-}) d\mathbf{y} ds$.

We highlight that this jump model generalizes discrete diffusion models discussed in Section 4.2, extending them from finite state spaces to continuous domains. Instead of continuous-time Markov chains in the finite-state setting, the forward and backward processes are now stochastic integrals with respect to Poisson random measures, and the score-matching loss is an integral instead of a sum. These differences pose computational challenges to the implementation of the model:

- **Forward Simulation:** By choosing a space and time-homogeneous Lévy measure $\lambda(\mathbf{y} - \mathbf{x})$, the forward process (5.3) reduces to a compound Poisson process. Specifically, in order to sample \mathbf{x}_t given \mathbf{x}_0 , one first samples the jump count from a Poisson distribution with mean $t \int_{\mathbb{R}^d} \lambda(\mathbf{y} - \mathbf{x}) d\mathbf{y}$, then samples corresponding jump locations from $\lambda(\cdot)$.
- **Numerical Integration:** If the Lévy measure $\lambda(\mathbf{y} - \mathbf{x}) d\mathbf{y}$ is further isotropic as in Example 4.4, the integral over $\lambda(\|\mathbf{y} - \mathbf{x}\|) d\mathbf{y}$ in the score-matching loss (5.4) is efficiently approximated by Monte Carlo integration using samples from the isotropic kernel $\lambda(\|\cdot\|)$.
- **Backward Simulation:** The backward simulation (5.5) is nontrivial due to the space-time inhomogeneity of the estimated Lévy measure $\tilde{\lambda}_t^\varphi(\mathbf{y}, \mathbf{x})$. Specifically, computing the integral:

$$\mathcal{J}[\tilde{\lambda}_t^\varphi(\cdot, \mathbf{y}_t^{(k)})] = \int_{\mathbb{R}^d} \tilde{\lambda}_t^\varphi(\mathbf{y}, \mathbf{y}_t^{(k)}) d\mathbf{y} = \int_{\mathbb{R}^d} \hat{s}_{T-t}^{\text{jump}}(\mathbf{y}_t^{(k)}, \mathbf{y}) \lambda(\|\mathbf{y} - \mathbf{y}_t^{(k)}\|) d\mathbf{y} \quad (5.6)$$

for each sample $\mathbf{y}_t^{(k)}$ at time t is computationally expensive. However, noting that the score function satisfies:

$$\hat{s}_t^{\text{jump}}(\mathbf{y}_t^{(k)}, \mathbf{y}) = \frac{\hat{s}_t^{\text{jump}}(\mathbf{y}^{\text{ref}}, \mathbf{y})}{\hat{s}_t^{\text{jump}}(\mathbf{y}^{\text{ref}}, \mathbf{y}_t^{(k)})},$$

with a fixed reference point \mathbf{y}^{ref} , one can reuse numerical integrations as

$$\mathcal{J}[\tilde{\lambda}_t^\varphi(\cdot, \mathbf{y}_t^{(k)})] = \frac{\mathcal{J}[\tilde{\lambda}_t^\varphi(\cdot, \mathbf{y}^{\text{ref}})]}{\hat{s}_{T-t}^{\text{jump}}(\mathbf{y}^{\text{ref}}, \mathbf{y}_t^{(k)})},$$

significantly reducing computational costs across samples. Subsequently, given the integral (5.6), one can sample the jump count from a Poisson distribution with mean $\kappa \mathcal{J}[\tilde{\lambda}_t^\varphi(\cdot, \mathbf{y}_t^{(k)})]$ and use rejection sampling to draw jump locations according to the estimated backward Lévy measure $\tilde{\lambda}_t^\varphi(\mathbf{y}, \mathbf{y}_t^{(k)})$, by first estimating $\sup_{\mathbf{y} \in \mathbb{R}^d} \hat{s}_t^{\text{jump}}(\mathbf{y}_t^{(k)}, \cdot)$ and then drawing samples from the kernel $\lambda(\|\mathbf{y} - \mathbf{y}_t^{(k)}\|)$.

In our empirical experiments, we consider a 2-D spatial domain defined on the torus $\mathbb{T}^2 = \mathbb{R}^2 / \mathbb{Z}^2$. The Lévy measure $\lambda_t(\mathbf{y}, \mathbf{x}) d\mathbf{y}$ is chosen as a time and space-homogeneous, isotropic Gaussian kernel:

$$\lambda_t(\mathbf{y}, \mathbf{x}) = \mathcal{N}(\mathbf{y} - \mathbf{x}; \mathbf{0}, \mathbf{I}) = \frac{1}{2\pi} \exp\left(-\frac{\|\mathbf{y} - \mathbf{x}\|^2}{2}\right).$$

We select a sufficiently large time horizon T such that the forward jump process converges towards the uniform distribution on \mathbb{T}^2 . To efficiently compute the integral term (5.6), we leverage the periodic structure of the domain and the property of the Gaussian kernel by employing numerical quadrature based on the fast Fourier transform. For all cases, we use an MLP with 5 layers and 128 hidden units to approximate the score function $\hat{s}_t^{\text{jump}}(\mathbf{x}, \mathbf{y})$.

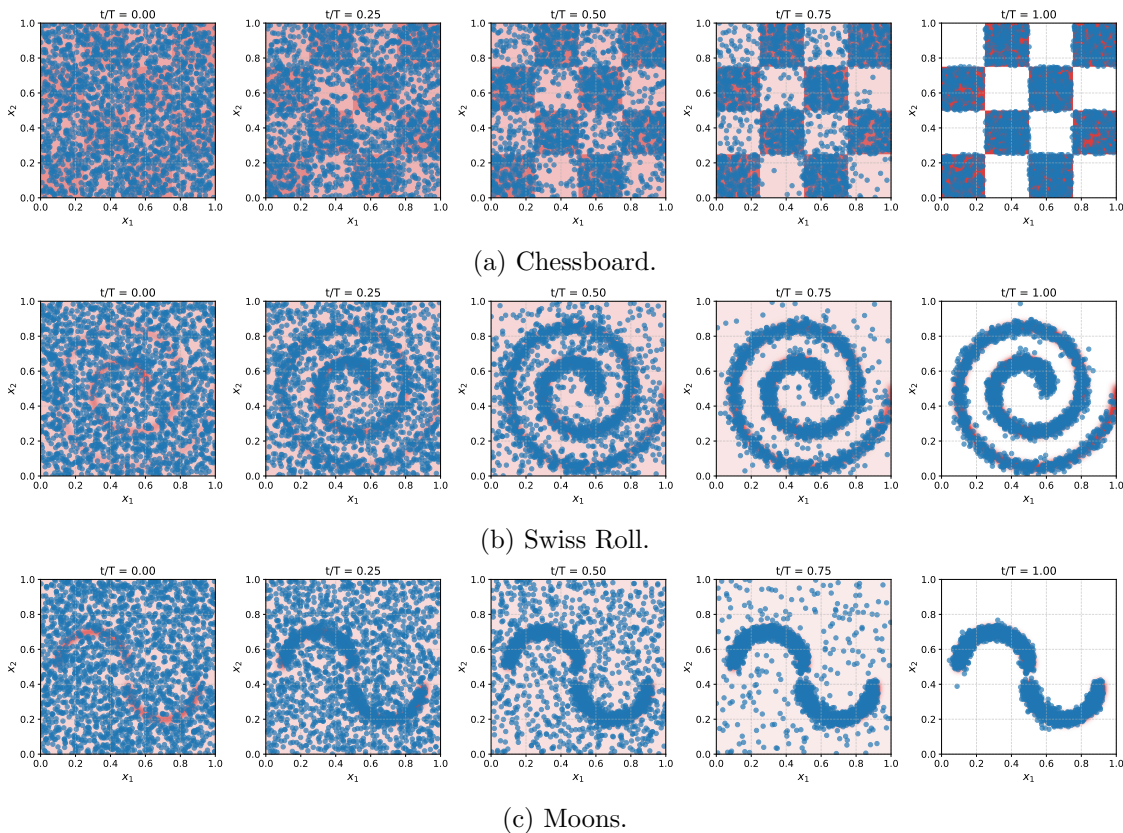


Figure 4: Visualization of generated backward trajectories using the jump process as the forward process. Samples from q_t generated via the learned backward dynamics (in blue) are compared with the corresponding distribution p_{T-t} from the forward process (in red). Results for three distributions (Chessboard, Swiss roll, Moons) demonstrate close agreement between the generated and target distributions.

We optimize the score-matching loss (5.4) using the Adam optimizer with a learning rate of 10^{-3} . Additionally, we use Fourier features for the input \mathbf{x} and \mathbf{y} to account for the periodic boundary conditions of the torus in the neural network architectures used to approximate the score functions.

The empirical results for three distinct 2-D target distributions (Chessboard, Swiss Roll, and Moons) are presented in Figure 4, illustrating the evolution of samples from the estimated backward jump dynamics. Specifically, we visualize 2048 backward-generated samples at different intermediate time steps and compare them with the corresponding forward distributions p_{T-t} . As demonstrated in Figure 4, the backward jump trajectories initially start from a uniform distribution at $t = 0$, gradually capturing intricate features of the target distributions as time progresses. By the end of the backward simulation ($t = T$), the generated samples accurately reconstruct the complex structures of each target distribution, verifying the effectiveness of our jump-based denoising Markov model.

Unlike diffusion models, which typically begin their backward evolution from Gaussian distributions, our jump-based model starts from the uniform distribution, highlighting a

distinct and versatile approach within generative modeling frameworks. We also refer readers to (Holderrieth et al., 2025, Figure 2 and Appendix E) for a comparison of our denoising approach with the generator matching framework in the context of pure jump processes. Our denoising Markov models do not directly work with the Kolmogorov’s forward equation to obtain the conditional generator $\tilde{\mathcal{L}}_{t|T}^{\tilde{\mathbf{x}}_T}$ through its coefficients $\tilde{\lambda}_{t|T}$, but instead leverage the score-matching technique to learn the backward process through the score function $\tilde{s}_t^{\text{jump}}$.

6. Discussion

In this paper, we introduced a novel framework for denoising Markov models, significantly extending classical diffusion models to a broader class of Markovian dynamics. By establishing connections between the backward process and the generalized Doob’s h -transform, we provided a rigorous theoretical foundation and a versatile methodology for constructing denoising Markov models from general Lévy-type forward processes. Under suitable regularity conditions, we justified the use of standard score-matching techniques for efficiently training these models. Additionally, we discussed practical implementation strategies and elaborated on connections with the generator matching approach recently proposed by Holderrieth et al. (2025). Through empirical experiments, we explored several novel instances of our general framework, including geometric Brownian motions and jump processes, highlighting their effectiveness in generative modeling scenarios that had not been previously studied. These experimental results demonstrated the flexibility and robustness of our approach, validating our theoretical insights.

For future work, several promising directions remain open. A key area of interest is extending the jump model to high-dimensional problems by introducing sparsity structures into the Lévy measure $\lambda_t(\mathbf{y}, \mathbf{x})$, following recent advances by Lou et al. (2024). This would enable efficient large-scale sampling and mass transportation, potentially unlocking powerful new generative modeling capabilities. Furthermore, a deeper theoretical investigation into how the choice of forward process affects training efficiency, generalization performance, and model robustness remains an exciting avenue. Specifically, it would be beneficial to systematically explore the optimal selection of forward processes tailored to specific structures present in target distributions, such as multimodality or heavy tails.

Acknowledgments

The authors thank the editor and the anonymous reviewers for their valuable comments and suggestions. Yinuo Ren and Lexing Ying were supported by the National Science Foundation under Grant No. DMS-2208163. This material is based upon work supported by the National Science Foundation under Grant No. CHE-2441297.

References

- 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.
- Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=li7qeBbCR1t>.
- Michael Samuel Albergo and Eric Vanden-Eijnden. Nets: A non-equilibrium transport sampler. In *International Conference on Machine Learning*, pages 1026–1055. PMLR, 2025.
- Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- David Applebaum. *Lévy processes and stochastic calculus*. Cambridge university press, 2009.
- Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021.
- Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. Multidiffusion: Fusing diffusion paths for controlled image generation. In *International Conference on Machine Learning*, pages 1737–1752. Pmlr, 2023.
- Joe Benton, Valentin De Bortoli, Arnaud Doucet, and George Deligiannidis. Nearly d -linear convergence bounds for diffusion models via stochastic localization. In *The Twelfth International Conference on Learning Representations*, 2024a. URL <https://openreview.net/forum?id=r5njV3BsuD>.
- Joe Benton, Yuyang Shi, Valentin De Bortoli, George Deligiannidis, and Arnaud Doucet. From denoising diffusions to denoising markov models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(2):286–301, 2024b.
- Adam Block, Youssef Mroueh, and Alexander Rakhlin. Generative modeling with denoising auto-encoders and langevin sampling. *arXiv preprint arXiv:2002.00107*, 2020.
- Björn Böttcher. Feller evolution systems: Generators and approximation. *Stochastics and Dynamics*, 14(03):1350025, 2014.
- Björn Böttcher, René Schilling, and Jian Wang. Lévy matters. iii. *Lecture Notes in Mathematics*, 2099:71–80, 2013.
- Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. *Advances in Neural Information Processing Systems*, 35:28266–28279, 2022.

- Patrick Cattiaux, Giovanni Conforti, Ivan Gentil, and Christian Léonard. Time reversal of diffusion processes under a finite entropy condition. In *Annales de l'Institut Henri Poincaré (B) Probabilités et Statistiques*, volume 59, pages 1844–1881. Institut Henri Poincaré, 2023.
- Haoxuan Chen, Yinuo Ren, Lexing Ying, and Grant M Rotskoff. Accelerating diffusion models with parallel sampling: Inference at sub-linear time complexity. *Advances in Neural Information Processing Systems*, 37:133661–133709, 2024a.
- Haoxuan Chen, Yinuo Ren, Martin Renqiang Min, Lexing Ying, and Zachary Izzo. Solving inverse problems via diffusion-based priors: An approximation-free ensemble sampling approach. *arXiv preprint arXiv:2506.03979*, 2025a.
- Hongrui Chen and Lexing Ying. Convergence analysis of discrete diffusion model: Exact implementation through uniformization. *Journal of Machine Learning*, 4(2):108–127, 2025.
- Hongrui Chen, Holden Lee, and Jianfeng Lu. Improved analysis of score-based generative modeling: User-friendly bounds under minimal smoothness assumptions. In *International Conference on Machine Learning*, pages 4735–4763. PMLR, 2023a.
- Junhua Chen, Lorenz Richter, Julius Berner, Denis Blessing, Gerhard Neumann, and Anima Anandkumar. Sequential controlled langevin diffusions. In *The Thirteenth International Conference on Learning Representations*, 2025b. URL <https://openreview.net/forum?id=dImD2sgy86>.
- Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=g7ohD1TITL>.
- Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *The Eleventh International Conference on Learning Representations*, 2023b. URL https://openreview.net/forum?id=zyLVMgsZOU_.
- Sitan Chen, Sinho Chewi, Holden Lee, Yuanzhi Li, Jianfeng Lu, and Adil Salim. The probability flow ode is provably fast. *Advances in Neural Information Processing Systems*, 36, 2024b.
- Xinlei Chen, Zhuang Liu, Saining Xie, and Kaiming He. Deconstructing denoising diffusion models for self-supervised learning. In *The Thirteenth International Conference on Learning Representations*, 2025c. URL <https://openreview.net/forum?id=9oMB6wnFYM>.
- Raphaël Chetrite and Shamik Gupta. Two refreshing views of fluctuation theorems through kinematics elements and exponential martingale. *Journal of Statistical Physics*, 143:543–584, 2011.
- Raphaël Chetrite and Hugo Touchette. Nonequilibrium markov processes conditioned on large deviations. In *Annales Henri Poincaré*, volume 16, pages 2005–2057. Springer, 2015.

- Muthu Chidambaram, Khashayar Gatmiry, Sitan Chen, Holden Lee, and Jianfeng Lu. What does guidance do? a fine-grained analysis in a simple setting. *Advances in Neural Information Processing Systems*, 37:84968–85005, 2024.
- Erhan Çinlar and Jean Jacod. Representation of semimartingale markov processes in terms of wiener processes and poisson random measures. In *Seminar on Stochastic Processes, 1981*, pages 159–242. Springer, 1981.
- Giovanni Conforti and Christian Léonard. Time reversal of markov processes with jumps under a finite entropy condition. *Stochastic Processes and their Applications*, 144:85–124, 2022.
- Philippe Courrège. Sur la forme intégro-différentielle des opérateurs de c_k^∞ dans c satisfaisant au principe du maximum. *Séminaire Brelot-Choquet-Deny. Théorie du Potentiel*, 10(1): 1–38, 1965.
- Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):10850–10869, 2023.
- Valentin De Bortoli, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. Riemannian score-based generative modelling. *Advances in Neural Information Processing Systems*, 35:2406–2422, 2022.
- Persi Diaconis and Laurent Miclo. On characterizations of metropolis type algorithms in continuous time. *ALEA: Latin American Journal of Probability and Mathematical Statistics*, 6:199–238, 2009.
- Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Genie: Higher-order denoising diffusion solvers. *Advances in Neural Information Processing Systems*, 35:30150–30166, 2022.
- Stewart N Ethier and Thomas G Kurtz. *Markov processes: characterization and convergence*. John Wiley & Sons, 2009.
- Wendell H Fleming and Halil Mete Soner. *Controlled Markov processes and viscosity solutions*, volume 25. Springer Science & Business Media, 2006.
- Nathan C. Frey, Dan Berenberg, Karina Zadorozhny, Joseph Kleinhenz, Julien Lafrance-Vanasse, Isidro Hotzel, Yan Wu, Stephen Ra, Richard Bonneau, Kyunghyun Cho, Andreas Loukas, Vladimir Gligorijevic, and Saeed Saremi. Protein discovery with discrete walk-jump sampling. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=zMPHK0mQNb>.
- Giovanni Gallavotti and Ezechiel Godert David Cohen. Dynamical ensembles in nonequilibrium statistical mechanics. *Physical review letters*, 74(14):2694, 1995.
- Xuefeng Gao and Lingjiong Zhu. Convergence analysis for general probability flow odes of diffusion models in wasserstein distances. In *International Conference on Artificial Intelligence and Statistics*, pages 1009–1017. PMLR, 2025.

- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Wei Guo, Molei Tao, and Yongxin Chen. Provable benefit of annealed langevin monte carlo for non-log-concave sampling. In *The Thirteenth International Conference on Learning Representations*, 2025a. URL <https://openreview.net/forum?id=P6IVIoGRRg>.
- Wei Guo, Molei Tao, and Yongxin Chen. Complexity analysis of normalizing constant estimation: from jarzynski equality to annealed importance sampling and beyond. *arXiv preprint arXiv:2502.04575*, 2025b.
- Zhiye Guo, Jian Liu, Yanli Wang, Mengrui Chen, Duolin Wang, Dong Xu, and Jianlin Cheng. Diffusion models in bioinformatics and computational biology. *Nature reviews bioengineering*, 2(2):136–154, 2024.
- Shivam Gupta, Linda Cai, and Sitan Chen. Faster diffusion sampling with randomized mid-points: Sequential and parallel. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=MT3a0fXIbY>.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47):1–33, 2022a.
- Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022b.
- Peter Holderrieth, Marton Havasi, Jason Yim, Neta Shaul, Itai Gat, Tommi Jaakkola, Brian Karrer, Ricky T. Q. Chen, and Yaron Lipman. Generator matching: Generative modeling with arbitrary markov processes. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=RuP17cJtZo>.
- Chin-Wei Huang, Milad Aghajohari, Joey Bose, Prakash Panangaden, and Aaron C Courville. Riemannian diffusion models. *Advances in Neural Information Processing Systems*, 35:2750–2761, 2022.
- Daniel Zhengyu Huang, Jiaoyang Huang, and Zhengjiang Lin. Convergence analysis of probability flow ode for score-based generative models. *IEEE Transactions on Information Theory*, 2025.
- Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.

- Christopher Jarzynski. Nonequilibrium equality for free energy differences. *Physical Review Letters*, 78(14):2690, 1997.
- Saravanan Kandasamy and Dheeraj Nagaraj. The poisson midpoint method for langevin dynamics: Provably efficient discretization for diffusion models. *Advances in Neural Information Processing Systems*, 37:65972–66024, 2024.
- Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Hiroshi Kunita. Absolute continuity of markov processes and generators. *Nagoya Mathematical Journal*, 36:1–26, 1969.
- Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence for score-based generative modeling with polynomial complexity. *Advances in Neural Information Processing Systems*, 35:22870–22882, 2022.
- Christian Léonard. Feynman-kac formula under a finite entropy condition. *Probability Theory and Related Fields*, 184(3):1029–1091, 2022.
- Gen Li and Yuchen Jiao. Improved convergence rate for diffusion probabilistic models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=S0d07Qxkw4>.
- Gen Li, Yu Huang, Timofey Efimov, Yuting Wei, Yuejie Chi, and Yuxin Chen. Accelerating convergence of score-based diffusion models, provably. In *International Conference on Machine Learning*, pages 27942–27954. PMLR, 2024a.
- Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. Towards non-asymptotic convergence for diffusion-based generative models. In *The Twelfth International Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=4VGEeER6W9>.
- Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems*, 35:4328–4343, 2022.
- 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. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.
- Xingchao Liu, Chengyue Gong, and qiang liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=XVjTT1nw5z>.
- Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=CNicRIVIPA>.

- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *Machine Intelligence Research*, 22(4):730–751, 2025.
- Nanye Ma, Mark Goldstein, Michael S Albergo, Nicholas M Boffi, Eric Vanden-Eijnden, and Saining Xie. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In *European Conference on Computer Vision*, pages 23–40. Springer, 2024.
- Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=aBsCjcPu_tE.
- Paul-André Meyer and Wei-An Zheng. Construction de processus de nelson réversibles. *Séminaire de probabilités de Strasbourg*, 19:12–26, 1985.
- Bernt Øksendal and Agnes Sulem. *Applied Stochastic Control of Jump Diffusions*. Springer, 2019.
- Zbigniew Palmowski and Tomasz Rolski. A technique for exponential change of measure for markov processes. *Bernoulli*, 8(6):767–785, 2002.
- Philip Protter. Point process differentials with evolving intensities. In *Nonlinear stochastic problems*, pages 467–472. Springer, 1983.
- Yinuo Ren, Haoxuan Chen, Grant M. Rotskoff, and Lexing Ying. How discrete and continuous diffusion meet: Comprehensive analysis of discrete diffusion models via a stochastic integral framework. In *The Thirteenth International Conference on Learning Representations*, 2025a. URL <https://openreview.net/forum?id=6awxwQEI82>.
- Yinuo Ren, Haoxuan Chen, Yuchen Zhu, Wei Guo, Yongxin Chen, Grant M. Rotskoff, Molei Tao, and Lexing Ying. Fast solvers for discrete diffusion models: Theory and applications of high-order algorithms. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025b. URL <https://openreview.net/forum?id=0uk1L6Q3s0>.
- Yinuo Ren, Wenhao Gao, Lexing Ying, Grant M. Rotskoff, and Jiequn Han. Driftlite: Lightweight drift control for inference-time scaling of diffusion models. In *The Fourteenth International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=101eG3Qik1>.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015.

- Lorenz Richter and Julius Berner. Improved sampling via learned diffusions. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=h4pNR0s006>.
- Nikil R Selvam, Amil Merchant, and Stefano Ermon. Self-refining diffusion samplers: Enabling parallelization via parareal iterations. *Advances in Neural Information Processing Systems*, 37:5429–5453, 2024.
- Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis Titsias. Simplified and generalized masked diffusion for discrete data. *Advances in neural information processing systems*, 37:103131–103167, 2024.
- Andy Shih, Suneel Belkhale, Stefano Ermon, Dorsa Sadigh, and Nima Anari. Parallel sampling of diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Marta Skreta, Tara Akhound-Sadegh, Viktor Ohanesian, Roberto Bondesan, Alan Aspuru-Guzik, Arnaud Doucet, Rob Brekelmans, Alexander Tong, and Kirill Neklyudov. Feynman-kac correctors in diffusion: Annealing, guidance, and product of experts. In *International Conference on Machine Learning*, pages 55906–55949. PMLR, 2025.
- 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.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in neural information processing systems*, 34: 1415–1428, 2021a.
- 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*, 2021b. URL <https://openreview.net/forum?id=PxtTIG12RRHS>.
- Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and Tommi Jaakkola. Dirichlet flow matching with applications to dna sequence design. In *International Conference on Machine Learning*, pages 46495–46513. PMLR, 2024.
- Belinda Tzen and Maxim Raginsky. Theoretical guarantees for sampling and inference in generative models with latent diffusions. In *Conference on Learning Theory*, pages 3084–3114. PMLR, 2019.
- Francisco Vargas, Will Sussman Grathwohl, and Arnaud Doucet. Denoising diffusion samplers. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=8pvnfTAbu1f>.

- Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- Yuqing Wang, Ye He, and Molei Tao. Evaluating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 37:19307–19352, 2024.
- Alexander D Wentzell, Hans Jürgen Engelbert, and Jürgen Groh. *Theorie zufälliger Prozesse*, volume 65. Birkhäuser Basel, Switzerland, 1979.
- Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7623–7633, 2023a.
- Tong Wu, Zhihao Fan, Xiao Liu, Hai-Tao Zheng, Yeyun Gong, Jian Jiao, Juntao Li, Jian Guo, Nan Duan, Weizhu Chen, et al. Ar-diffusion: Auto-regressive diffusion model for text generation. *Advances in Neural Information Processing Systems*, 36:39957–39974, 2023b.
- Yuchen Wu, Yuxin Chen, and Yuting Wei. Stochastic runge-kutta methods: Provable acceleration of diffusion models. *arXiv preprint arXiv:2410.04760*, 2024.
- Zhen Xing, Qijun Feng, Haoran Chen, Qi Dai, Han Hu, Hang Xu, Zuxuan Wu, and Yugang Jiang. A survey on video diffusion models. *ACM Computing Surveys*, 57(2):1–42, 2024.
- Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=PzcvxEMzvQC>.
- Minkai Xu, Tomas Geffner, Karsten Kreis, Weili Nie, Yilun Xu, Jure Leskovec, Stefano Ermon, and Arash Vahdat. Energy-based diffusion language models for text generation. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=sL2F9YCMXf>.
- Ruihan Yang, Prakhar Srivastava, and Stephan Mandt. Diffusion probabilistic modeling for video generation. *Entropy*, 25(10):1469, 2023.
- Sherry Yang, KwangHwan Cho, Amil Merchant, Pieter Abbeel, Dale Schuurmans, Igor Mordatch, and Ekin Dogus Cubuk. Scalable diffusion for materials generation. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=wm4WlHoXpC>.
- Eun Bi Yoon, Keehun Park, Sungwoong Kim, and Sungbin Lim. Score-based generative models with lévy processes. *Advances in Neural Information Processing Systems*, 36:40694–40707, 2023.

- Linfeng Zhang, Weinan E, and Lei Wang. Monge-ampère flow for generative modeling. *arXiv preprint arXiv:1809.10188*, 2018.
- Qinsheng Zhang and Yongxin Chen. Path integral sampler: A stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=_uCb2ynRu7Y.
- Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=Loek7hfb46P>.
- Kaiwen Zheng, Cheng Lu, Jianfei Chen, and Jun Zhu. Dpm-solver-v3: Improved diffusion ode solver with empirical model statistics. *Advances in Neural Information Processing Systems*, 36:55502–55542, 2023.
- Yuchen Zhu, Tianrong Chen, Ling kai Kong, Evangelos Theodorou, and Molei Tao. Trivialized momentum facilitates diffusion generative modeling on lie groups. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=DTatjJTD11>.

Appendix A. Mathematical Approach to Denoising Markov Models

In this section, we assume that E is a locally compact, separable Hausdorff space, equipped with a base Radon measure $(E, \mathcal{B}(E), \mu)$, where $\mathcal{B}(E)$ denotes the Borel σ -algebra on E . Let $(x_t)_{t \in \mathbb{T}}$ be a time-inhomogeneous Markov process on E , defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where \mathbb{T} is the time index set, either \mathbb{R} or $[0, +\infty)$. For notational simplicity, we will write $\mathbb{P}(\cdot)$ instead of $\mathbb{P}(\{\omega \in \Omega \mid \cdot\})$ when there is no risk of confusion.

A.1 Feller Evolution Systems

We begin by introducing the concept of time-evolution operators and evolution systems, which form a family of linear operators closely associated with Markov processes.

Definition A.1 (Time-Evolution Operators) *A family of linear operators $(U_{t,s})_{s \leq t, s, t \in \mathbb{T}}$ is called a time-evolution operator family, or simply an evolution system, if the following conditions hold:*

- (1) $U_{t,t} = \text{id}$ for any $t \in \mathbb{T}$;
- (2) $U_{t,s} = U_{t,r}U_{r,s}$ for any $s \leq r \leq t$, with $s, r, t \in \mathbb{T}$.

This definition is motivated by interpreting the operator $U_{t,s}$ as the conditional expectation operator associated with the Markov process $(x_t)_{t \in \mathbb{T}}$. Specifically, for any $f \in B_b(E)$, the space of bounded Borel measurable functions on E , and any $s \leq t$, we define the operator $U_{t,s}$ by

$$U_{t,s}f(x) = \mathbb{E}[f(x_t) \mid x_s = x]. \quad (\text{A.1})$$

It is straightforward to verify that $U_{t,s}$ satisfies both conditions in Definition A.1, and hence qualifies as an evolution system. Moreover, the operator family $(U_{t,s})$ defined in (A.1) possesses the following properties:

- (1) (Positivity-Preserving Property) $U_{t,s}f \geq 0$, if $f \geq 0$ and $f \in B_b(E)$, for any $s \leq t$, $s, t \in \mathbb{T}$;
- (2) (Contractivity) $\|U_{t,s}f\|_\infty \leq \|f\|_\infty$, for any $f \in B_b(E)$, $s \leq t$, $s, t \in \mathbb{T}$;
- (3) (Conservation) $U_{t,s}1 = 1$, for any $s \leq t$, $s, t \in \mathbb{T}$.

We also define a special class of evolution systems, namely the *Feller evolution system* (Böttcher, 2014), as follows:

Definition A.2 (Feller Evolution System) *An evolution system $(U_{t,s})_{s \leq t, s, t \in \mathbb{T}}$ is called a Feller evolution system if, for any $f \in C_0(E)$, the space of continuous functions vanishing at infinity¹, the following conditions are satisfied:*

- (1) (Feller Property) $U_{t,s}f \in C_0(E)$, for any $s \leq t$, $s, t \in \mathbb{T}$;
- (2) (Positivity-Preserving Property) $U_{t,s}f \geq 0$ if $f \geq 0$, for any $s \leq t$, $s, t \in \mathbb{T}$;

1. A function f is said to *vanish at infinity* if for any $\epsilon > 0$, there exists a compact set $K \subset E$ such that $|f(x)| < \epsilon$ for all $x \in E \setminus K$.

(3) (Contractivity) $\|U_{t,s}f\|_\infty \leq \|f\|_\infty$, for any $s \leq t$, $s, t \in \mathbb{T}$;

(4) (Strong Continuity) The following limit holds:

$$\lim_{(\sigma,\tau) \rightarrow (s,t)} \|U_{\sigma,\tau}f - U_{t,s}f\|_\infty = 0, \text{ for any } s \leq t, s, t \in \mathbb{T}.$$

For any evolution system, we can define its right and left generators as follows.

Definition A.3 (Generator of Evolution System) For any fixed $s \in \mathbb{R}$, the right generator of the evolution system is defined by the limit

$$\mathcal{L}_s f = \lim_{h \rightarrow 0^+} \frac{U_{s+h,s}f - f}{h}.$$

The domain $\text{dom}(\mathcal{L}_s)$ of the right generator \mathcal{L}_s at time s consists of all functions $f \in C_b(E)$ such that the above limit exists in the L^∞ norm, in which $C_b(E)$ denotes the space of bounded continuous functions on E .

Similarly, the left generator is defined as

$$\mathcal{L}_s^- f = \lim_{h \rightarrow 0^+} \frac{U_{s,s-h}f - f}{h},$$

with the corresponding domain denoted by $\text{dom}(\mathcal{L}_s^-)$.

A.2 Feller Semigroups

For time-homogeneous evolution systems, observe that for any $s \in \mathbb{T}$ and $t \geq 0$, we have

$$T_t := U_{s+t,s} = U_{s,s-t},$$

which implies that the left and right generators coincide, *i.e.*, $\mathcal{L} = \mathcal{L}_s = \mathcal{L}_s^-$. In this case, the family of linear operators $(T_t)_{t \geq 0}$ forms a semigroup, or more precisely, a *one-parameter semigroup*, as defined below:

Definition A.4 (One-Parameter Semigroup) A family of linear operators $(T_t)_{t \geq 0}$ is called a one-parameter semigroup if the following conditions are satisfied:

- (1) $T_0 = \text{id}$;
- (2) $T_s T_t = T_{s+t}$, for any $s, t \geq 0$.

For such semigroups, we define their generator as follows:

Definition A.5 (Generator of Semigroup) The generator of a semigroup is given by the limit

$$\mathcal{L}f = \lim_{h \rightarrow 0} \frac{T_h f - f}{h},$$

where the domain $\text{dom}(\mathcal{L})$ consists of all functions $f \in C_b(E)$ for which the above limit exists in the L^∞ norm.

In general, the generator \mathcal{L} is a closed operator on its domain, which is dense in $C_b(E)$, and it maps $\text{dom}(\mathcal{L})$ into $C_b(E)$ (Böttcher et al., 2013). To focus on a class of semigroups that are particularly suitable for Markovian dynamics, we now introduce the notion of a Feller semigroup.

Definition A.6 (Feller Semigroup) *A one-parameter semigroup $(T_t)_{t \geq 0}$ is called a Feller semigroup if the following conditions hold for every $f \in C_0(E)$:*

- (1) (Feller Property) $T_t f \in C_0(E)$ for any $t \geq 0$;
- (2) (Positivity-Preserving Property) $T_t f \geq 0$ if $f \geq 0$, for any $t \geq 0$;
- (3) (Contractivity) $\|T_t f\|_\infty \leq \|f\|_\infty$ for any $t \geq 0$;
- (4) (Strong Continuity) The following limit holds:

$$\lim_{h \rightarrow 0^+} \|T_h f - f\|_\infty = 0.$$

We now introduce the concept of a Feller process, which is a time-homogeneous Markov process whose transition semigroup forms a Feller semigroup.

Definition A.7 (Feller Process) *A time-homogeneous Markov process $(x_t)_{t \in \mathbb{T}}$ on E is called a Feller process if its transition semigroup, defined as*

$$T_t f(x) := \mathbb{E}[f(x_t) | x_0 = x]$$

is a Feller semigroup.

An important identity associated with Feller processes and their generators is Dynkin's formula. This result provides a key connection between the infinitesimal generator and the expected value of a function evaluated along the process.

Theorem A.8 (Dynkin's Formula (Böttcher et al., 2013)) *For any $f \in \text{dom}(\mathcal{L}) \cap C_0(E)$ and stopping time t with $\mathbb{E}[t | x_0 = x] < +\infty$, the following relation holds:*

$$\mathbb{E}[f(x_t) | x_0 = x] - f(x_0) = \mathbb{E} \left[\int_0^t \mathcal{L} f(x_s) ds \middle| x_0 = x \right].$$

To establish a more general framework that connects Feller evolution systems with Feller semigroups, we introduce an augmentation construction. This construction allows us to reinterpret a time-inhomogeneous process as a time-homogeneous one by lifting it into an augmented state space.

Definition A.9 (Augmented Process (Wentzell et al., 1979)) *Let $(x_t)_{t \in \mathbb{T}}$ be a Markov process on E with the evolution system $(U_{t,s})_{s \leq t, s, t \in \mathbb{T}}$. Then we define the following augmented process $(\tilde{x}_s)_{s \geq 0}$ in the augmented probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ defined as follows:*

- (1) (State Space) The augmented state \tilde{x} is defined as $\tilde{x} := (s, x) \in \mathbb{T} \times E$, with the augmented state space $\mathbb{T} \times E$ equipped with the following σ -algebra:

$$\tilde{\mathcal{B}} := \left\{ \tilde{B} \subset \mathbb{T} \times E \mid \tilde{B}_s := \left\{ x \in E \mid \tilde{x} = (s, x) \in \tilde{B} \right\} \in \mathcal{B}(E), \forall s \in \mathbb{T} \right\}$$

where $\mathcal{B}(E)$ is the Borel σ -algebra on E ;

- (2) (Sample Space) The augmented event $\tilde{\omega}$ is defined as $\tilde{\omega} := (s, \omega) \in \mathbb{T} \times \Omega := \tilde{\Omega}$, with the augmented sample space $\tilde{\Omega}$ equipped with the following σ -algebra:

$$\tilde{\mathcal{F}} := \left\{ \tilde{F} \subset \tilde{\Omega} \mid \tilde{F}_s := \left\{ \omega \in \Omega \mid \tilde{\omega} = (s, \omega) \in \tilde{F} \right\} \in \mathcal{F}, \forall s \in \mathbb{T} \right\};$$

- (3) (Probability Measure) The augmented process $(\tilde{x}_t)_{t \geq 0}$ is defined as

$$\tilde{x}_t(\tilde{\omega}) = \tilde{x}_t(s, \omega) := (s + t, x_{s+t}(\omega)),$$

and the augmented probability measure $\tilde{\mathbb{P}}$ is defined such that the following relation holds:

$$\begin{aligned} \tilde{\mathbb{P}}_{\tilde{x}_0} \left(\tilde{x}_t(\tilde{\omega}) \in \tilde{B} \right) &:= \tilde{\mathbb{P}} \left(\tilde{x}_t(\tilde{\omega}) \in \tilde{B} \mid \tilde{x}_0(\tilde{\omega}) = \tilde{x} \right) \\ &= \tilde{\mathbb{P}} \left((s + t, x_{s+t}(\omega)) \in \tilde{B} \mid \tilde{x}_0(\omega) = (s, x_s(\omega)) = (s, x) \right) \\ &:= \mathbb{P} \left(x_{s+t}(\omega) \in \tilde{B}_{s+t} \mid x_s(\omega) = x \right), \end{aligned}$$

for any $\tilde{x} = (s, x) \in \mathbb{T} \times E$, $t \geq 0$, and $\tilde{B} \in \tilde{\mathcal{B}}$.

For simplicity, we will omit the explicit dependence on ω and $\tilde{\omega}$ in the remainder of our discussion.

The following proposition establishes that the augmented process is indeed a Markov process on $\mathbb{T} \times E$, provided that the original process is Markovian on E .

Proposition A.10 *If the original process $(x_t)_{t \in \mathbb{T}}$ is a Markov process on E , then the augmented process $(\tilde{x}_t)_{t \in \mathbb{T}}$ is a Markov process on $\mathbb{T} \times E$ with the following transition probability:*

$$\tilde{\mathbb{P}} \left(\tilde{x}_{t'} \in \tilde{B} \mid \tilde{x}_t = \tilde{x} \right) = \mathbb{P} \left(x_{s'+t'-t} \in \tilde{B}_{s'+t'-t} \mid x_{s'} = x \right),$$

for any $\tilde{x} = (s', x) \in \mathbb{T} \times E$, $\tilde{B} \in \tilde{\mathcal{B}}$, and $s', t' \geq t$.

Proof Let $\tilde{F} \in \tilde{\mathcal{F}}_t$ and $\tilde{B} \in \tilde{\mathcal{B}}$ with $t' \geq t$, where $\tilde{\mathcal{F}}_t$ denotes the filtration generated by the augmented process up to time t . We verify the Markov property of the augmented process by computing:

$$\begin{aligned} &\int_{\tilde{F}} \mathbf{1}_{\tilde{B}}(\tilde{x}_{t'}) \tilde{\mathbb{P}}_{\tilde{x}_0} (d\tilde{\omega}) = \int_{\tilde{F}} \mathbf{1}_{\tilde{B}}(\tilde{x}_{t'}) \tilde{\mathbb{P}} (d\tilde{\omega} \mid \tilde{x}_0 = (s, x)) \\ &= \int_{\tilde{F}_s} \mathbf{1}_{\tilde{B}_{s+t'}}(x_{s+t'}) \mathbb{P}(d\omega \mid x_s = x) \\ &= \int_{\tilde{F}_s} \mathbb{P} \left(x_{s+t'} \in \tilde{B}_{s+t'} \mid x_{s+t} \right) \mathbb{P}(d\omega \mid x_s = x) \\ &= \int_{\tilde{F}} \mathbb{P} \left(\tilde{x}_{t'} \in \tilde{B} \mid \tilde{x}_t \right) \tilde{\mathbb{P}} (d\tilde{\omega} \mid \tilde{x}_0 = \tilde{x}), \end{aligned}$$

where the second equality is due to that conditioning $\tilde{x}_0(\tilde{\omega}) = (s, x)$ implies that $\tilde{\omega} = (s, \omega)$ such that $x_s(\omega) = x$, and the third equality is by the Markov property of the original process $(x_t)_{t \in \mathbb{T}}$.

We now compute the transition probability of the augmented process explicitly. For any $x \in E$, $\tilde{B} \in \tilde{\mathcal{B}}$, and $s', t' \geq t$, we have

$$\begin{aligned} & \tilde{\mathbb{P}}\left(\tilde{x}_{t'} \in \tilde{B} \mid \tilde{x}_t = \tilde{x}\right) \\ &= \tilde{\mathbb{P}}\left((s' + t' - t, x_{s'+t'-t}) \in \tilde{B} \mid \tilde{x}_t = (s', x)\right) \\ &= \mathbb{P}\left(x_{s'+t'-t} \in \tilde{B}_{s'+t'-t} \mid x_{s'} = x\right). \end{aligned}$$

This shows that the transition kernel depends only on the difference $t' - t$, and hence the augmented process is time-homogeneous. \blacksquare

We may now define the semigroup associated with the augmented time-homogeneous Markov process:

$$\tilde{T}_t f(\tilde{x}) = \tilde{\mathbb{E}}[f(\tilde{x}_t) \mid \tilde{x}_0 = \tilde{x}],$$

where $\tilde{\mathbb{E}}$ denotes the expectation with respect to the augmented probability measure $\tilde{\mathbb{P}}$. Correspondingly, the generator of the augmented process is introduced below.

Definition A.11 (Augmented Generator) *The augmented generator $\tilde{\mathcal{L}}$ of the augmented process is defined as the following limit:*

$$\tilde{\mathcal{L}}f(\tilde{x}) = \lim_{h \rightarrow 0} \frac{\tilde{T}_h f(\tilde{x}) - f(\tilde{x})}{h},$$

for any $f \in C_b(\mathbb{T} \times E)$. We denote the domain $\text{dom}(\tilde{\mathcal{L}})$ of the augmented generator $\tilde{\mathcal{L}}$ as the space of all function $f \in C_b(\mathbb{T} \times E)$ satisfying the conditions aforementioned such that the limit above exists with respect to L^∞ norm.

The next proposition establishes the relation between the augmented generator and the right generator of the original time-inhomogeneous Markov process.

Proposition A.12 *Let $(x_t)_{t \in \mathbb{T}}$ be a Markov process governed by the evolution system $(U_{t,s})_{s \leq t}$ with right generator \mathcal{L}_s . Let $(\tilde{x}_t)_{t \geq 0}$ be the corresponding augmented process with semigroup $(\tilde{T}_t)_{t \geq 0}$ and generator $\tilde{\mathcal{L}}$ as in Definition A.11.*

Suppose $f \in C_b([0, T] \times E)$ satisfies the following conditions:

- (1) $f(\cdot, x) \in C^1(\mathbb{T})$ for any $x \in E$;
- (2) $f(s, \cdot), \partial_s f(s, \cdot) \in \text{dom}(\mathcal{L}_s)$ for any $s \in \mathbb{T}$,

then we have $f \in \text{dom}(\tilde{\mathcal{L}})$ and the generator satisfies

$$\tilde{\mathcal{L}}f(\tilde{x}) = \tilde{\mathcal{L}}f(s, x) = \partial_s f(s, x) + \mathcal{L}_s f(s, x). \quad (\text{A.2})$$

Proof Starting from the definition, we write

$$\begin{aligned}
 \tilde{\mathcal{L}}f(s, x) &= \lim_{h \rightarrow 0} \frac{\tilde{T}_h f(s, x) - f(s, x)}{h} \\
 &= \lim_{h \rightarrow 0} \frac{\mathbb{E}[f(s+h, x_{s+h})|x_s = x] - f(s, x)}{h} \\
 &= \lim_{h \rightarrow 0} \frac{\mathbb{E}[f(s+h, x_{s+h}) - f(s, x_{s+h}) + f(s, x_{s+h}) - f(s, x)|x_s = x]}{h} \\
 &= \lim_{h \rightarrow 0} \frac{\mathbb{E}[f(s+h, x_{s+h}) - f(s, x_{s+h})|x_s = x]}{h} + \lim_{h \rightarrow 0} \frac{\mathbb{E}[f(s, x_{s+h}) - f(s, x)|x_s = x]}{h},
 \end{aligned}$$

where the second term on the right-hand side matches the definition of the right generator \mathcal{L}_s of the original process $(x_t)_{t \in \mathbb{T}}$ and thus converges to $\mathcal{L}_s f(s, x)$ as $h \rightarrow 0$, given $f(s, \cdot) \in \text{dom}(\mathcal{L}_s)$.

The first term on the right-hand side can be computed as follows:

$$\begin{aligned}
 &\left| \frac{1}{h} \mathbb{E}[f(s+h, x_{s+h}) - f(s, x_{s+h})|x_s = x] - \partial_s f(s, x_s) \right| \\
 &\leq \left| \frac{1}{h} \mathbb{E}[f(s+h, x_{s+h}) - f(s, x_{s+h})|x_s = x] - \mathbb{E}[\partial_s f(s, x_{s+h})|x_s = x] \right| \\
 &\quad + \left| \mathbb{E}[\partial_s f(s, x_{s+h})|x_s = x] - \partial_s f(s, x) \right| \\
 &= \left| \frac{1}{h} \mathbb{E} \left[\int_s^{s+h} (\partial_\sigma f(\sigma, x_{s+h}) - \partial_s f(s, x_{s+h})) d\sigma \middle| x_s = x \right] \right| \\
 &\quad + \left| \mathbb{E}[\partial_s f(s, x_{s+h})|x_s = x] - \partial_s f(s, x) \right|,
 \end{aligned}$$

where the first term is bounded by the continuity of $\partial_s f(\cdot, x)$ as

$$\begin{aligned}
 &\left| \frac{1}{h} \mathbb{E} \left[\int_s^{s+h} (\partial_\sigma f(\sigma, x_{s+h}) - \partial_s f(s, x_{s+h})) d\sigma \middle| x_s = x \right] \right| \\
 &\leq \sup_{\sigma \in [s, s+h], x \in E} |\partial_\sigma f(\sigma, x) - \partial_s f(s, x)| \rightarrow 0, \text{ as } h \rightarrow 0,
 \end{aligned}$$

and the second term is bounded as

$$\begin{aligned}
 &\left| \mathbb{E}[\partial_s f(s, x_{s+h})|x_s = x] - \partial_s f(s, x) \right| \\
 &= U_{s+h, s} \partial_s f(s, \cdot)(x) - \partial_s f(s, x) \rightarrow 0, \text{ as } h \rightarrow 0,
 \end{aligned}$$

where we treat the derivative $\partial_s f(s, x)$ as a function of only x and thus the limit converges to 0 as $h \rightarrow 0$, given $\partial_s f(s, \cdot) \in \text{dom}(\mathcal{L}_s)$. The proof is thus completed. \blacksquare

We now state an important corollary, which extends Dynkin's formula to time-inhomogeneous processes via the augmented framework. We will adopt the shorthand notation $f_t(\cdot) = f(t, \cdot)$ in the following discussions.

Corollary A.13 (Dynkin's Formula) *Under the conditions of Proposition A.12, the following relation holds:*

$$\mathbb{E}[f_t(x_t)|x_0 = x] - f_0(x) = \mathbb{E} \left[\int_0^t (\partial_s f_s(x_s) + \mathcal{L}_s f_s(x_s)) ds \right],$$

for any stopping time t with $\mathbb{E}[t|x_0 = x] < +\infty$.

Proof First, we apply Theorem A.8 to the augmented process $(\tilde{x}_t)_{t \in \mathbb{T}}$, and obtain the following relation:

$$\tilde{\mathbb{E}}[f(\tilde{x}_t)|\tilde{x}_0 = (0, x)] - f(0, x) = \tilde{\mathbb{E}} \left[\int_0^t \tilde{\mathcal{L}}f(\tilde{x}_s) ds \middle| \tilde{x}_0 = (0, x) \right],$$

where the left-hand side, by Definition A.9, can be rewritten as

$$\tilde{\mathbb{E}}[f(\tilde{x}_t)|\tilde{x}_0 = (0, x)] - f(0, x) = \mathbb{E}[f_t(x_t)|x_0 = x] - f_0(x),$$

while the right-hand side, similarly, can be rewritten as

$$\tilde{\mathbb{E}} \left[\int_0^t \tilde{\mathcal{L}}f(\tilde{x}_s) ds \middle| \tilde{x}_0 = (0, x) \right] = \mathbb{E} \left[\int_0^t (\partial_s f_s(x_s) + \mathcal{L}_s f_s(x_s)) ds \right],$$

where we used the relation (A.2) in the last equality. The proof is thus completed. \blacksquare

Finally, the next theorem establishes the equivalence between Feller evolution systems and Feller semigroups via the augmented process.

Theorem A.14 (Böttcher (2014, Theorem 3.2)) *Let $(x_t)_{t \in \mathbb{T}}$ be a Markov process on E and $(\tilde{x}_t)_{t \in \mathbb{T}}$ be the augmented process on $\mathbb{T} \times E$ as defined above. Then the following statements are equivalent:*

- (1) *The evolution system $(U_{t,s})_{s \leq t}$ is a Feller evolution system;*
- (2) *The semigroup $(\tilde{T}_t)_{t \in \mathbb{T}}$ is a Feller semigroup, and thus the augmented process $(\tilde{x}_t)_{t \in \mathbb{T}}$ is a Feller process.*

Remark A.15 *The time index set \mathbb{T} may also be taken as a closed interval $[0, T]$ for some $T > 0$. In this case, we may restrict all time indices in the preceding discussion to lie within $[0, T]$. The identity in Eq. (A.2) remains valid for all $s \in [0, T)$ since it is a local property. To ensure well-defined behavior at the endpoint T , one may extend the Markov process $(x_t)_{t \in [0, T]}$ to $(x_t)_{t \in [0, T + \epsilon]}$ for some small $\epsilon > 0$, where $T^+ = T + \epsilon$, or more conveniently, to $(x_t)_{t \in \mathbb{R}}$.*

A.3 Time-Reversal of Markov Processes

Throughout this subsection, we fix the time index set as $\mathbb{T} = [0, T]$ and adopt the following assumption.

Assumption A.16 (Feller) *The Markov process $(x_t)_{t \in [0, T]}$ is governed by a Feller evolution system $(U_{t,s})_{s \leq t, s, t \in [0, T]}$. Therefore, by Theorem A.14, the associated augmented process $(\tilde{x}_t)_{t \in [0, T]}$ is a Feller process.*

Additionally, we assume that both the original and augmented processes are càdlàg.

We now define the time reversal of a Markov process, following the formulation in Léonard (2022); Cattiaux et al. (2023).

Definition A.17 (Time-Reversal of Markov Process) *Let $(x_t)_{t \in [0, T]}$ be a Markov process. The associated time-reversal process $(\tilde{x}_t)_{t \in [0, T]}$ is defined by*

$$\tilde{x}_t := \lim_{h \rightarrow 0^+} x_{T-t-h} = x_{(T-t)^-},$$

and by setting $\tilde{x}_T := x_0$.

Without loss of generality, we assume that $x_{T-} = x_T$ with probability 1. Consequently, the map $\tilde{\cdot} : (x_t)_{t \in [0, T]} \mapsto (\tilde{x}_t)_{t \in [0, T]}$ is almost surely bijective. The filtration $(\tilde{\mathcal{F}}_t)_{t \in [0, T]}$ generated by the time-reversal process is defined by

$$\tilde{\mathcal{F}}_t = \bigcap_{s > t} \sigma(x_{T-s:T}),$$

where $\sigma(x_{s:T})$ denotes the σ -algebra generated by the path $(x_t)_{t \in [s, T]}$. We define the right generator $\tilde{\mathcal{L}}_s$ of the time-reversal process $(\tilde{x}_t)_{t \in [0, T]}$ and the generator $\tilde{\tilde{\mathcal{L}}}_s$ of the augmented time-reversal process $(\tilde{\tilde{x}}_t)_{t \in [0, T]}$, similar to Definition A.3 and Definition A.5.

For the ease of presentation, we introduce the following function class assumption.

Assumption A.18 (Function Class \mathcal{U}) *There exists a class of functions \mathcal{U} on E satisfying the following:*

- (1) \mathcal{U} is an algebra, i.e., closed under addition, (pointwise) multiplication, and scalar multiplication;
- (2) $\mathcal{U} \subset C_b(E)$ and is dense in $C_b(E)$;
- (3) $\mathcal{U} \subset \bigcap_{t \in [0, T]} \text{dom}(\mathcal{L}_t) \cap \text{dom}(\tilde{\mathcal{L}}_t)$.

Next, we define the formal adjoint of a generator. This concept is motivated by the observation that we only require adjointness to hold for functions in \mathcal{U} , rather than for all elements of $L^2(E, \mu)$.

Definition A.19 (Adjoint Operator) *We define the domain $\text{dom}(\mathcal{L}^*)$ of the adjoint operator \mathcal{L}^* as the space of all functions $g \in L^2(E, \mu)$ such that there exists a unique function $\mathcal{L}^*g \in L^2(E, \mu)$ satisfying that for any $f \in \mathcal{U}$, the following relation holds:*

$$\langle \mathcal{L}f, g \rangle = \langle f, \mathcal{L}^*g \rangle, \tag{A.3}$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product on $L^2(E, \mu)$.

The uniqueness of \mathcal{L}^*g follows from the density of \mathcal{U} in $C_b(E)$ and hence in $L^2(E, \mu)$. We now introduce another function class $\mathcal{V} \subset \mathcal{U}$ as test functions for the function class \mathcal{U} :

Assumption A.20 (Function Class \mathcal{V}) *There exists a class of functions \mathcal{V} on E satisfying:*

- (1) $\mathcal{V} \subset \bigcap_{t \in [0, T]} \text{dom}(\mathcal{L}_t^*)$;
- (2) $g\mathcal{U} = \{fg | f \in \mathcal{U}\} \in \mathcal{V}$, for any $g \in \mathcal{V}$.

To ensure well-definedness and integrability of expressions involving the generator, we also make the following regularity assumption on marginal and conditional densities.

Assumption A.21 (Regularity of Densities) *We assume that*

- (1) *The marginal distribution of the Markov process $(x_t)_{t \in [0, T]}$ at time t admits a density $p_t(x)$ with respect to the base measure μ and $p_t \in \mathcal{V} \subset \mathcal{U}$, for any $t \in [0, T]$;*
- (2) *The conditional distribution $p_{t|s}(x_t|x_s)$ of the Markov process $(x_t)_{t \in [0, T]}$, defined as*

$$p_{t|s}(\cdot|x_s) = \mathbb{P}(x_t \in \cdot | x_s = x_s),$$

admits a density $p_{t|s}(x_t|x_s)$ with respect to the base measure μ and $p_{t|s}(\cdot|x_s) \in \mathcal{V} \subset \mathcal{U}$ for any $x_s \in E$, for any $s, t \in [0, T]$ with $s < t$;

- (3) *For any $f \in \mathcal{U}$, the following conditions hold:*

$$\mathbb{E} \left[\int_0^T |\mathcal{L}_t f(x_t)| dt \right] < +\infty, \quad \text{and} \quad \mathbb{E} \left[\int_0^T |\tilde{\mathcal{L}}_t f(\tilde{x}_t)| dt \right] < +\infty.$$

We are now ready to state a key result that relates the generator of the original process to that of the time-reversal process.

Theorem A.22 (Time-Reversal Generator) *Under Assumptions A.18, A.20 and A.21, for each $t \in [0, T]$ and any $f \in \mathcal{U}$, the generator $\tilde{\mathcal{L}}_t$ of the time-reversal process satisfies the identity*

$$p_t \tilde{\mathcal{L}}_{T-t} f = \mathcal{L}_t^*(p_t f) - f \mathcal{L}_t^* p_t. \tag{A.4}$$

If in addition $f \in \mathcal{V}$ and $p_t(x) > 0$ for all $x \in E$, the identity Eq. (A.4) can be rewritten as

$$\tilde{\mathcal{L}}_{T-t} f = p_t^{-1} \mathcal{L}_t^*(p_t f) - p_t^{-1} f \mathcal{L}_t^* p_t = \mathcal{L}_t^* f + p_t^{-1} \Gamma_t^*(p_t, f)$$

where $\Gamma_t^(p_t, f) = \mathcal{L}_t^*(p_t f) - p_t \mathcal{L}_t^* f - f \mathcal{L}_t^* p_t$ is the carré du champ operator associated with the adjoint operator \mathcal{L}_t^* .*

Proof By Theorem 3.17 in Cattiaux et al. (2023), for any $t \in [0, T]$ and $f, g \in \mathcal{U}$, the following relation holds:

$$\mathbb{E}_{x_t \sim p_t} \left[(\mathcal{L}_t + \tilde{\mathcal{L}}_{T-t}) f(x_t) g(x_t) + \Gamma_t(f, g)(x_t) \right] = 0,$$

where $\Gamma_t(f, g) = \mathcal{L}_t(fg) - f \mathcal{L}_t g - g \mathcal{L}_t f$ is the carré du champ operator of the original generator \mathcal{L}_t .

Substituting the definition of $\Gamma_t(f, g)$, we obtain

$$\begin{aligned} \langle p_t \tilde{\mathcal{L}}_{T-t} f, g \rangle &= -\langle p_t \mathcal{L}_t f, g \rangle - \langle p_t, \Gamma_t(f, g) \rangle \\ &= -\langle p_t \mathcal{L}_t f, g \rangle - \langle p_t, \mathcal{L}_t(fg) - f\mathcal{L}_t g - g\mathcal{L}_t f \rangle \\ &= \langle p_t f, \mathcal{L}_t g \rangle - \langle p_t, \mathcal{L}_t(fg) \rangle. \end{aligned}$$

Since $p_t, f \in \mathcal{V}$ and $fg \in \mathcal{U}$, the definition of the adjoint operator yields

$$\begin{aligned} \langle p_t \tilde{\mathcal{L}}_{T-t} f, g \rangle &= \langle p_t f, \mathcal{L}_t g \rangle - \langle p_t, \mathcal{L}_t(fg) \rangle \\ &= \langle \mathcal{L}_t^*(p_t f), g \rangle - \langle f \mathcal{L}_t^* p_t, g \rangle \\ &= \langle \mathcal{L}_t^*(p_t f) - f \mathcal{L}_t^* p_t, g \rangle, \end{aligned}$$

which establishes the identity (A.4). ■

For completeness, we state the Kolmogorov forward equation under the current framework.

Theorem A.23 (Kolmogorov Forward Equation) *Under Assumptions A.18, A.20 and A.21, the transition density $p_{t|0}(x|x_0)$ satisfies the evolution equation*

$$\partial_t p_{t|0}(x|x_0) = \mathcal{L}_t^* p_{t|0}(x|x_0).$$

In particular, the marginal density satisfies

$$\partial_t p_t = \mathcal{L}_t^* p_t.$$

Proof Applying Dynkin's formula (Corollary A.13), for any $f \in \mathcal{U}$, we compute

$$\begin{aligned} \langle f, \partial_t p_{t|0}(\cdot|x_0) \rangle &= \partial_t \langle f, p_{t|0}(\cdot|x_0) \rangle = \partial_t \mathbb{E}[f(x_t)|x_0 = x] \\ &= \mathbb{E}[(\partial_t + \mathcal{L}_t)f(x_t)|x_0 = x] = \mathbb{E}[\mathcal{L}_t f(x_t)|x_0 = x], \end{aligned}$$

which can also be written as

$$\mathbb{E}[\mathcal{L}_t f(x_t)|x_0 = x] = \langle \mathcal{L}_t f, p_{t|0}(\cdot|x_0) \rangle = \langle f, \mathcal{L}_t^* p_{t|0}(\cdot|x_0) \rangle,$$

and the proof is thus complete. ■

A.4 Change of Measure

Motivated by the form of the time-reversal generator $\tilde{\mathcal{L}}_{T-t}$ in Theorem A.22, we now consider the following structural assumption for a family of generators $(\mathcal{K}_t)_{t \in [0, T]}$ associated with another Markov process $(y_t)_{t \in [0, T]}$. This formulation provides a rigorous version of Assumption 3.3 in the main text.

Assumption A.24 (Parametrization of Backward Generator) *There exists a family of generators $(\mathcal{K}_t)_{t \in [0, T]}$ corresponding to a Markov process $(y_t)_{t \in [0, T]}$ such that, for each $t \in [0, T]$, the generator \mathcal{K}_t and the adjoint operator \mathcal{L}_t^* are related by a function $\varphi_t \in \mathcal{U}$ in the following way:*

$$\varphi_t \mathcal{K}_{T-t} f = \mathcal{L}_t^*(\varphi_t f) - f \mathcal{L}_t^* \varphi_t, \quad (\text{A.5})$$

for all $f \in \mathcal{U}$.

To ensure that the identity in (A.5) is well-defined, we require the following regularity properties of the density ratio $\eta_t := \varphi_t/p_t$, which informally ensures that this ratio is smooth and bounded.

Assumption A.25 (Regularity of η_t) *For each $t \in [0, T]$, the function $\eta_t := \varphi_t p_t^{-1}$ satisfies:*

(1) $\eta(\cdot), \eta^{-1}(\cdot) \in C_b([0, T] \times E)$, $\eta(x), \eta(x)^{-1} \in C^1([0, T])$ for any $x \in E$, and

$$\eta_t, \partial_t \eta_t, \eta_t^{-1}, \partial_t \eta_t^{-1}, \log \eta_t, \partial_t \log \eta_t \in \mathcal{U}$$

for any $t \in [0, T]$;

(2) The function $\tilde{\mathcal{L}}_{T-t} \eta_t$, satisfying that

$$\tilde{\mathcal{L}}_{T-t} \eta_t = p_t^{-1} \mathcal{L}_t^*(p_t \eta_t) - p_t^{-1} \eta_t \mathcal{L}_t^* p_t,$$

is upper bounded, i.e.,

$$\sup_{x \in E} \tilde{\mathcal{L}}_{T-t} \eta_t(x) = \sup_{x \in E} \frac{p_t(x) \mathcal{L}_t^* \varphi_t(x) - \varphi_t(x) \mathcal{L}_t^* p_t(x)}{p_t^2(x)} < +\infty.$$

We note that condition (2) in Assumption A.25 implicitly requires both p_t and φ_t to be strictly positive on E . Under Assumption A.25 (1), it is easy to see that $\varphi_t = p_t \eta_t \in \mathcal{V}$ and also $\varphi_t f = p_t \eta_t f \in \mathcal{V}$ for any $f \in \mathcal{U}$, and thus Eq. (A.5) is well-defined under Definition A.19, i.e., $\mathcal{U} \subset \text{dom}(\mathcal{K}_t)$ for any $t \in [0, T]$. Analogous to Theorem A.22, if we further assume that $f \in \mathcal{V}$ and $\varphi_t(x) > 0$ for any $x \in E$, Eq. (A.5) can be rewritten as

$$\mathcal{K}_{T-t} f = \varphi_t^{-1} \mathcal{L}_t^*(\varphi_t f) - \varphi_t^{-1} f \mathcal{L}_t^* \varphi_t = \mathcal{L}_t^* f + \varphi_t^{-1} \Gamma_t^*(\varphi_t, f).$$

We now express the generator \mathcal{K}_t in terms of the time-reversal generator $\tilde{\mathcal{L}}_{T-t}$ and the ratio $\eta_t = \varphi_t/p_t$.

Lemma A.26 *For any $t \in [0, T]$, under Assumptions A.24 and A.25, the following relation holds for any $f \in \mathcal{U}$:*

$$\eta_t \mathcal{K}_{T-t} f = \tilde{\mathcal{L}}_{T-t}(\eta_t f) - f \tilde{\mathcal{L}}_{T-t} \eta_t = \eta_t \tilde{\mathcal{L}}_{T-t} f + \tilde{\Gamma}_{T-t}(\eta_t, f),$$

where $\tilde{\Gamma}_t(f, g) = \tilde{\mathcal{L}}_t(fg) - f \tilde{\mathcal{L}}_t g - g \tilde{\mathcal{L}}_t f$ is the carré du champ operator associated with the time-reversal generator $\tilde{\mathcal{L}}_t$.

Proof For any functions f, g that allow the following computation (cf. Assumptions A.18 and A.20), using the definition of η_t , we have that

$$\left\langle g, \tilde{\mathcal{L}}_{T-t}(\eta_t f) - f \tilde{\mathcal{L}}_{T-t} \eta_t \right\rangle = \left\langle gp_t^{-1}, p_t \tilde{\mathcal{L}}_{T-t}(\varphi_t p_t^{-1} f) \right\rangle - \left\langle gp_t^{-1}, f p_t \tilde{\mathcal{L}}_{T-t}(\varphi_t p_t^{-1}) \right\rangle.$$

Notice that by the time-reversal formula (Theorem A.22), we have

$$p_t \tilde{\mathcal{L}}_{T-t}(\varphi_t p_t^{-1} f) = \mathcal{L}_t^*(\varphi_t f) - \varphi_t p_t^{-1} f \mathcal{L}_t^* p_t \quad \text{and} \quad p_t \tilde{\mathcal{L}}_{T-t}(\varphi_t p_t^{-1}) = \mathcal{L}_t^*(\varphi_t) - \varphi_t p_t^{-1} \mathcal{L}_t^* p_t,$$

and a careful algebraic manipulation shows that

$$\begin{aligned} & \left\langle g, \tilde{\mathcal{L}}_{T-t}(\eta_t f) - f \tilde{\mathcal{L}}_{T-t} \eta_t \right\rangle \\ &= \left\langle gp_t^{-1}, \mathcal{L}_t^*(\varphi_t f) \right\rangle - \left\langle gp_t^{-1}, \varphi_t p_t^{-1} f \mathcal{L}_t^* p_t \right\rangle - \left\langle gp_t^{-1}, f \mathcal{L}_t^* \varphi_t \right\rangle + \left\langle gp_t^{-1}, f \varphi_t p_t^{-1} \mathcal{L}_t^* p_t \right\rangle \\ &= \left\langle gp_t^{-1}, \mathcal{L}_t^*(\varphi_t f) \right\rangle - \left\langle gp_t^{-1}, f \mathcal{L}_t^* \varphi_t \right\rangle = \left\langle gp_t^{-1}, \mathcal{L}_t^*(\varphi_t f) - f \mathcal{L}_t^* \varphi_t \right\rangle, \end{aligned}$$

which by Assumption A.24 is exactly

$$\left\langle g, \tilde{\mathcal{L}}_{T-t}(\eta_t f) - f \tilde{\mathcal{L}}_{T-t} \eta_t \right\rangle = \left\langle gp_t^{-1}, \varphi_t \mathcal{K}_{T-t} f \right\rangle = \left\langle g, \eta_t \mathcal{K}_{T-t} f \right\rangle,$$

and the proof is complete. \blacksquare

Corollary A.27 *Under Assumptions A.24 and A.25, suppose any $f \in C_b([0, T] \times E)$ satisfying that*

- (1) $f(\cdot, x) \in C^1([0, T])$ for any $x \in E$;
- (2) $f(s, \cdot), \partial_s f(s, \cdot) \in \mathcal{U}$ for any $s \in \mathbb{T}$,

then the following relation holds:

$$\eta_{T-}(\cdot) \tilde{\mathcal{K}} f = \tilde{\mathcal{L}}(\eta_{T-}(\cdot) f) - f \tilde{\mathcal{L}} \eta_{T-}(\cdot) = \eta_{T-}(\cdot) \tilde{\mathcal{L}} f + \tilde{\Gamma}(\eta_{T-}(\cdot), f), \quad (\text{A.6})$$

where $\tilde{\mathcal{K}}$ is the augmented generator of the augmented process $(\tilde{y}_t)_{t \in [0, T]}$ defined according to Definition A.9, and $\tilde{\Gamma}(f, g) = \tilde{\mathcal{L}}(fg) - f \tilde{\mathcal{L}} g - g \tilde{\mathcal{L}} f$ is the carré du champ operator associated with the augmented time-reversal generator $\tilde{\mathcal{L}}$.

Proof We first apply Proposition A.12 to the augmented time-reversal process $(\tilde{x}_t)_{t \in [0, T]}$ to obtain that the augmented generator $\tilde{\mathcal{L}}_t$ satisfies that

$$\tilde{\mathcal{L}} f(s, x) = \partial_s f(s, x) + \tilde{\mathcal{L}}_s f(s, x), \quad (\text{A.7})$$

and then apply similar arguments to the Markov process $(y_t)_{t \in [0, T]}$ assumed in Assumption A.24, we also have that

$$\tilde{\mathcal{K}} f(s, x) = \partial_s f(s, x) + \mathcal{K}_s f(s, x).$$

Using Lemma A.26, we compute

$$\begin{aligned}
 \eta_{T-s}(x)\tilde{\mathcal{K}}f(s,x) &= \eta_{T-s}(x)\partial_s f(s,x) + \eta_{T-s}(x)\mathcal{K}_s f(s,x) \\
 &= \eta_{T-s}(x)\partial_s f(s,x) + \tilde{\mathcal{L}}_s(\eta_{T-s}f(s,\cdot))(x) - f(s,x)\tilde{\mathcal{L}}_s\eta_{T-s}(x) \\
 &= \partial_s(\eta_{T-s}(x)f(s,x)) + \tilde{\mathcal{L}}_s(\eta_{T-s}f(s,\cdot))(x) - f(s,x)\left(\partial_s\eta_{T-s}(x) + \tilde{\mathcal{L}}_s\eta_{T-s}(x)\right) \\
 &= \tilde{\mathcal{L}}(\eta_{T-\cdot}(\cdot)f) - f\tilde{\mathcal{L}}(\eta_{T-\cdot}(\cdot)),
 \end{aligned}$$

which proves the result. \blacksquare

Corollary A.28 *Under the conditions of Corollary A.27, we further have that*

$$\eta_{T-\cdot}^{-1}(\cdot)\tilde{\mathcal{L}}f = \tilde{\mathcal{K}}(\eta_{T-\cdot}^{-1}(\cdot)f) - f\tilde{\mathcal{K}}\eta_{T-\cdot}^{-1}(\cdot) = \eta_{T-\cdot}^{-1}(\cdot)\tilde{\mathcal{K}}f + \tilde{\Gamma}_{\mathcal{K}}(\eta_{T-\cdot}^{-1}(\cdot), f),$$

where $\tilde{\Gamma}_{\mathcal{K}}(f,g) = \tilde{\mathcal{K}}(fg) - f\tilde{\mathcal{K}}g - g\tilde{\mathcal{K}}f$ is the carré du champ operator associated with the augmented generator \mathcal{K} .

Proof By taking $f = \eta_{T-\cdot}^{-1}(\cdot)$ in Eq. (A.6), we have

$$\eta_{T-\cdot}(\cdot)\tilde{\mathcal{K}}\eta_{T-\cdot}^{-1}(\cdot) = \tilde{\mathcal{L}}1 - \eta_{T-\cdot}^{-1}(\cdot)\tilde{\mathcal{L}}\eta_{T-\cdot}(\cdot) = -\eta_{T-\cdot}^{-1}(\cdot)\tilde{\mathcal{L}}\eta_{T-\cdot}(\cdot),$$

and also by taking $f = \eta_{T-\cdot}^{-1}(\cdot)g$ in Eq. (A.6), we have

$$\begin{aligned}
 \eta_{T-\cdot}(\cdot)\tilde{\mathcal{K}}(\eta_{T-\cdot}^{-1}(\cdot)g) &= \tilde{\mathcal{L}}(\eta_{T-\cdot}(\cdot)\eta_{T-\cdot}^{-1}(\cdot)g) - g\eta_{T-\cdot}^{-1}(\cdot)\tilde{\mathcal{L}}\eta_{T-\cdot}(\cdot) \\
 &= \tilde{\mathcal{L}}g + g\eta_{T-\cdot}(\cdot)\tilde{\mathcal{K}}\eta_{T-\cdot}^{-1}(\cdot),
 \end{aligned}$$

and thus

$$\eta_{T-\cdot}^{-1}(\cdot)\tilde{\mathcal{L}}g = \tilde{\mathcal{K}}(\eta_{T-\cdot}^{-1}(\cdot)g) - g\tilde{\mathcal{K}}\eta_{T-\cdot}^{-1}(\cdot),$$

and the proof is complete. \blacksquare

The relation of \mathcal{K}_t (resp., $\tilde{\mathcal{K}}$) and the time-reversal generator $\tilde{\mathcal{L}}_t$ (resp., $\tilde{\mathcal{L}}$) in Lemma A.26 can be regarded as applying to $\tilde{\mathcal{L}}_t$ (resp., $\tilde{\mathcal{L}}$) a specific form of “perturbation” involving the carré du champ operator acting with the ratio η_{T-t} :

$$\eta_{T-t}^{-1}\tilde{\Gamma}_t(\eta_{T-t}, \cdot) \quad (\text{resp.} \quad \eta_{T-\cdot}(\cdot)^{-1}\tilde{\Gamma}(\eta_{T-\cdot}(\cdot), \cdot)),$$

which allows an exponential change of measure argument, as summarized in the following theorem:

Theorem A.29 (Change of Measure) *Let \mathbb{Q} be another probability measure absolutely continuous with respect to \mathbb{P} . Suppose that, for any $t \in [0, T]$, the conditional expected log-Radon-Nikodym derivative satisfies*

$$\mathbb{E} \left[\log \frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{\tilde{\mathcal{F}}_t} \right] = \mathbb{E} \left[- \int_0^t \left(\eta_{T-s}^{-1}(\tilde{x}_s)\tilde{\mathcal{L}}_s\eta_{T-s}(\tilde{x}_s) - \tilde{\mathcal{L}}_s \log \eta_{T-s}(\tilde{x}_s) \right) ds \right], \quad (\text{A.8})$$

then the time-reversal process $(\tilde{x}_t)_{t \in [0, T]}$ under the probability measure \mathbb{Q} is governed by the generator \mathcal{K}_t .

Furthermore, the original probability measure \mathbb{P} is also absolutely continuous to \mathbb{Q} , and the conditional expected log-Radon-Nikodym derivative above satisfies the following relation:

$$\mathbb{E} \left[\log \frac{d\mathbb{P}}{d\mathbb{Q}} \Big|_{\tilde{\mathcal{F}}_t} \right] = \mathbb{E} \left[\int_0^t \left(\eta_{T-s}^{-1}(\tilde{x}_s) \tilde{\mathcal{L}}_s \eta_{T-s}(\tilde{x}_s) - \tilde{\mathcal{L}}_s \log \eta_{T-s}(\tilde{x}_s) \right) ds \right].$$

Proof Based on Corollary A.27, we apply Theorem 4.2 in Palmowski and Rolski (2002) to the augmented generator $\tilde{\mathcal{L}}$ of the augmented time-reversal process $(\tilde{x}_t)_{t \in [0, T]}$ by verifying condition (M2) in Proposition 3.2 therein, and obtain that under another probability measure $\tilde{\mathbb{Q}}$ absolutely continuous to $\tilde{\mathbb{P}}$, the augmented time-reversal process $(\tilde{x}_t)_{t \in [0, T]}$ is governed by the augmented generator $\tilde{\mathcal{K}}$ and the Radon-Nikodym derivative is given by

$$\frac{d\tilde{\mathbb{Q}}}{d\tilde{\mathbb{P}}}\Big|_{\tilde{\mathcal{F}}_t} = \frac{\eta_{T-t}(\tilde{x}_t)}{\eta_T(\tilde{x}_0)} \exp \left(- \int_0^t \eta_{T-s}^{-1}(\tilde{x}_s) \tilde{\mathcal{L}} \eta_{T-s}(\tilde{x}_s) ds \right),$$

and thus

$$\begin{aligned} \mathbb{E} \left[\log \frac{d\tilde{\mathbb{Q}}}{d\tilde{\mathbb{P}}}\Big|_{\tilde{\mathcal{F}}_t} \right] &= \mathbb{E} \left[- \int_0^t \eta_{T-s}^{-1}(\tilde{x}_s) \tilde{\mathcal{L}} \eta_{T-s}(\tilde{x}_s) ds + \log \eta_{T-t}(\tilde{x}_t) - \log \eta_T(\tilde{x}_0) \right] \\ &= \mathbb{E} \left[- \int_0^t \left(\eta_{T-s}^{-1}(\tilde{x}_s) \tilde{\mathcal{L}} \eta_{T-s}(\tilde{x}_s) - \tilde{\mathcal{L}} \log \eta_{T-s}(\tilde{x}_s) \right) ds \right], \end{aligned}$$

where the boundary terms $\log \eta_{T-t}(\tilde{x}_t) - \log \eta_T(\tilde{x}_0)$ are rewritten via Dynkin's formula (Theorem A.8) applied to the time-inhomogeneous function $(s, x) \mapsto \log \eta_{T-s}(x)$ on the interval $[0, t]$.

Notice that in the construction of the augmented processes (Definition A.9), no extra randomness is introduced apart from that from the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and thus we may define the probability measure \mathbb{Q} on the original probability space $(\Omega, \mathcal{F}, \mathbb{P})$ under which the time-reversal process $(\tilde{x}_t)_{t \in [0, T]}$ is governed by the generator \mathcal{K}_t and the Radon-Nikodym derivative is given by

$$\begin{aligned} \mathbb{E} \left[\log \frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{\tilde{\mathcal{F}}_t} \right] &:= \mathbb{E} \left[\log \frac{d\tilde{\mathbb{Q}}}{d\tilde{\mathbb{P}}}\Big|_{\tilde{\mathcal{F}}_t} \right] \\ &= \mathbb{E} \left[- \int_0^t \left(\eta_{T-s}^{-1}(\tilde{x}_s) \tilde{\mathcal{L}} \eta_{T-s}(\tilde{x}_s) - \tilde{\mathcal{L}} \log \eta_{T-s}(\tilde{x}_s) \right) ds \right] \\ &= \mathbb{E} \left[- \int_0^t \left(\eta_{T-s}^{-1}(\tilde{x}_s) \left(\partial_s \eta_{T-s}(\tilde{x}_s) + \tilde{\mathcal{L}}_s \eta_{T-s}(\tilde{x}_s) \right) - \left(\partial_s \log \eta_{T-s}(\tilde{x}_s) + \tilde{\mathcal{L}}_s \log \eta_{T-s}(\tilde{x}_s) \right) \right) ds \right] \\ &= \mathbb{E} \left[- \int_0^t \left(\eta_{T-s}^{-1}(\tilde{x}_s) \tilde{\mathcal{L}}_s \eta_{T-s}(\tilde{x}_s) - \tilde{\mathcal{L}}_s \log \eta_{T-s}(\tilde{x}_s) \right) ds \right], \end{aligned}$$

where the second equality is due to Eq. (A.7) and the last equality is by

$$\eta_{T-s} \partial_s \log \eta_{T-s} = \partial_s \eta_{T-s}.$$

To prove the absolute continuity of \mathbb{P} to \mathbb{Q} , one applies the same argument to the augmented generator $\tilde{\mathcal{K}}$ and the original generator \mathcal{K}_t of the Markov process $(y_t)_{t \in [0, T]}$ under the probability measure \mathbb{Q} and $\tilde{\mathbb{Q}}$, respectively, given the symmetry of the two augmented generators $\tilde{\mathcal{K}}$ and $\tilde{\mathcal{L}}$ (Corollary A.28), and the proof is thus completed. \blacksquare

Proof of Corollary 3.7 The error bound is obtained from Theorem 3.6 by

$$D_{\text{KL}}(p_0 \| q_T) \leq D_{\text{KL}}(p_{0:T} \| q_{0:T}) = D_{\text{KL}}(p_T \| q_0) + D_{\text{KL}}(p_{0:T} \| q_{0:T} | p_T), \quad (\text{A.9})$$

where the first inequality is by applying the data-processing inequality and the second equality is by the chain rule of KL divergence. The conditional KL divergence $D_{\text{KL}}(p_{0:T} \| q_{0:T} | p_T)$, defined as

$$D_{\text{KL}}(p_{0:T} \| q_{0:T} | p_T) = \mathbb{E}_{\tilde{x}_0 \sim p_T} [D_{\text{KL}}(p_{0:T}(\tilde{x}_{0:T}) \| q_{0:T}(\tilde{x}_{0:T})) | \tilde{x}_0],$$

can be further bounded by the change-of-measure theorem Theorem 3.6 by aligning the two processes $(\tilde{x}_t)_{t \in [0, T]}$ and $(y_t)_{t \in [0, T]}$ at time 0 with probability distribution $\tilde{p}_0 = p_T$. \blacksquare

Corollary A.30 *Under the conditions in Theorem A.29, the Radon-Nikodym derivative (A.8) further satisfies the following relation:*

$$D_{\text{KL}}(\mathbb{P} \| \mathbb{Q}) = \mathbb{E} \left[\log \frac{d\mathbb{P}}{d\mathbb{Q}} \right] = \mathbb{E} \left[\int_0^T \eta_s(x_s) \mathcal{L}_s \eta_s^{-1}(x_s) + \mathcal{L}_s \log \eta_s(x_s) ds \right] := \mathfrak{L}[\eta_t].$$

Proof By Theorem A.29, we have that

$$\begin{aligned} D_{\text{KL}}(\mathbb{P} \| \mathbb{Q}) &= \mathbb{E} \left[\log \frac{d\mathbb{P}}{d\mathbb{Q}} \right] \\ &= \mathbb{E} \left[\int_0^T \eta_{T-s}^{-1}(\tilde{x}_s) \tilde{\mathcal{L}}_s \eta_{T-s}(\tilde{x}_s) - \tilde{\mathcal{L}}_s \log \eta_{T-s}(\tilde{x}_s) ds \right] \\ &= \int_0^T \mathbb{E}_{\tilde{x}_s \sim p_{T-s}} \left[\eta_{T-s}^{-1}(\tilde{x}_s) \tilde{\mathcal{L}}_s \eta_{T-s}(\tilde{x}_s) - \tilde{\mathcal{L}}_s \log \eta_{T-s}(\tilde{x}_s) \right] ds \\ &= \int_0^T \left(\langle p_{T-s}, \eta_{T-s}^{-1} \tilde{\mathcal{L}}_s \eta_{T-s} \rangle - \langle p_{T-s}, \tilde{\mathcal{L}}_s \log \eta_{T-s} \rangle \right) ds \\ &= \int_0^T \left(\langle \eta_{T-s}^{-1}, p_{T-s} \tilde{\mathcal{L}}_s \eta_{T-s} \rangle - \langle 1, p_{T-s} \tilde{\mathcal{L}}_s \log \eta_{T-s} \rangle \right) ds. \end{aligned}$$

Notice that by Theorem A.22, we have the following time-reversal formula:

$$p_{T-s} \tilde{\mathcal{L}}_s \eta_{T-s} = \mathcal{L}_{T-s}^*(p_{T-s} \eta_{T-s}) - \eta_{T-s} \mathcal{L}_{T-s}^* p_{T-s},$$

and

$$p_{T-s} \tilde{\mathcal{L}}_s \log \eta_{T-s} = \mathcal{L}_{T-s}^*(p_{T-s} \log \eta_{T-s}) - \log \eta_{T-s} \mathcal{L}_{T-s}^* p_{T-s}.$$

Thus, we continue our computation as follows:

$$\begin{aligned}
 & D_{\text{KL}}(\mathbb{P} \parallel \mathbb{Q}) \\
 &= \int_0^T \langle \eta_{T-s}^{-1}, \mathcal{L}_{T-s}^*(p_{T-s} \eta_{T-s}) - \eta_{T-s} \mathcal{L}_{T-s}^* p_{T-s} \rangle ds \\
 &\quad - \int_0^T \langle 1, \mathcal{L}_{T-s}^*(p_{T-s} \log \eta_{T-s}) - \log \eta_{T-s} \mathcal{L}_{T-s}^* p_{T-s} \rangle ds \\
 &= \int_0^T \langle \eta_{T-s}^{-1}, \mathcal{L}_{T-s}^*(p_{T-s} \eta_{T-s}) \rangle - \langle \eta_{T-s}^{-1}, \eta_{T-s} \mathcal{L}_{T-s}^* p_{T-s} \rangle ds \\
 &\quad - \int_0^T \langle 1, \mathcal{L}_{T-s}^*(p_{T-s} \log \eta_{T-s}) \rangle - \langle 1, \log \eta_{T-s} \mathcal{L}_{T-s}^* p_{T-s} \rangle ds \\
 &= \int_0^T \langle \mathcal{L}_{T-s} \eta_{T-s}^{-1}, p_{T-s} \eta_{T-s} \rangle - \langle 1, \mathcal{L}_{T-s}^* p_{T-s} \rangle ds \\
 &\quad - \int_0^T \langle \mathcal{L}_{T-s} 1, p_{T-s} \log \eta_{T-s} \rangle - \langle \log \eta_{T-s}, \mathcal{L}_{T-s}^* p_{T-s} \rangle ds \\
 &= \int_0^T (\langle \eta_{T-s} \mathcal{L}_{T-s} \eta_{T-s}^{-1}, p_{T-s} \rangle + \langle \mathcal{L}_{T-s} \log \eta_{T-s}, p_{T-s} \rangle) ds \\
 &= \int_0^T (\langle \eta_s \mathcal{L}_s \eta_s^{-1}, p_s \rangle + \langle \mathcal{L}_s \log \eta_s, p_s \rangle) ds \\
 &= \int_0^T \mathbb{E}_{x_s \sim p_s} [\eta_s(x_s) \mathcal{L}_s \eta_s^{-1}(x_s) + \mathcal{L}_s \log \eta_s(x_s)] ds
 \end{aligned}$$

where the third equality is by noticing that $\mathcal{L}_t 1 = 0$ by the definition of the generator. The proof is thus complete. \blacksquare

Remark A.31 (Direct route via generalized Doob's h -transform) *The conclusion above may also be obtained more directly from Theorem 3.5. Indeed, by Lemma 3.4, after replacing t by $T - t$, we have*

$$\mathcal{K}_t f = \eta_{T-t}^{-1} \tilde{\mathcal{L}}_t (\eta_{T-t} f) - \eta_{T-t}^{-1} f \tilde{\mathcal{L}}_t \eta_{T-t}.$$

Therefore \mathcal{K}_t is precisely the generalized Doob's h -transform of the true backward generator $\tilde{\mathcal{L}}_t$ with

$$h_t = \eta_{T-t}, \quad \lambda_t = \eta_{T-t}^{-1} \tilde{\mathcal{L}}_t \eta_{T-t}.$$

Applying Theorem 3.5 to the time-reversal process $(\tilde{x}_t)_{t \in [0, T]}$ yields

$$\frac{d\mathbb{Q}}{d\mathbb{P}}(\tilde{x}_{[0, T]}) = \frac{\eta_0(\tilde{x}_T)}{\eta_T(\tilde{x}_0)} \exp\left(-\int_0^T \left(\eta_{T-t}^{-1} \partial_t \eta_{T-t} + \eta_{T-t}^{-1} \tilde{\mathcal{L}}_t \eta_{T-t}\right)(\tilde{x}_t) dt\right).$$

Taking logarithms, inverting the Radon–Nikodym derivative, and then applying Dynkin's formula to the time-inhomogeneous function $(t, x) \mapsto \log \eta_{T-t}(x)$ give

$$D_{\text{KL}}(\mathbb{P} \parallel \mathbb{Q}) = \mathbb{E}_{\mathbb{P}} \left[\int_0^T \left(\eta_{T-t}^{-1} \tilde{\mathcal{L}}_t \eta_{T-t} - \tilde{\mathcal{L}}_t \log \eta_{T-t} \right) (\tilde{x}_t) dt \right].$$

Finally, Theorem 3.2 rewrites this in terms of the forward generator as

$$D_{\text{KL}}(\mathbb{P}||\mathbb{Q}) = \mathbb{E}_{\mathbb{P}} \left[\int_0^T (\eta_t \mathcal{L}_t \eta_t^{-1} + \mathcal{L}_t \log \eta_t)(x_t) dt \right],$$

which is exactly the identity in Corollary A.30.

The next corollary rewrites the Radon-Nikodym derivative (A.8) in terms of the conditional distribution $p_{t|0}$.

Corollary A.32 (Score-Matching) *Suppose the discussions above can be extended to the case where $p_0 = \delta_{x_0}$ for any $x_0 \in E$, then there exists a constant C such that*

$$D_{\text{KL}}(\mathbb{P}||\mathbb{Q}) = \mathbb{E} \left[\log \frac{d\mathbb{P}}{d\mathbb{Q}} \right] = \mathfrak{L}_{\text{SM}}[\eta_{t|0}] + C,$$

where we define the score-matching loss $\mathfrak{L}_{\text{SM}}[\eta_{t|0}]$ as

$$\mathfrak{L}_{\text{SM}}[\eta_{t|0}] = \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T \mathbb{E}_{x_s \sim p_{s|0}(\cdot|x_0)} \left[\eta_{s|0}(x_s|x_0) \mathcal{L}_s \eta_{s|0}^{-1}(x_s|x_0) + \mathcal{L}_s \log \eta_{s|0}(x_s|x_0) \right] ds \right],$$

the constant C only depends on the Markov process x_t , or more precisely, the marginal distribution p_t , the initial distribution p_0 , and the conditional distributions $p_{t|0}(\cdot|x_0)$, and does not depend on φ_t .

Proof Using similar arguments as in the proof of Corollary A.30, we have

$$\begin{aligned} D_{\text{KL}}(\mathbb{P}||\mathbb{Q}) &= \mathbb{E} \left[\log \frac{d\mathbb{P}}{d\mathbb{Q}} \right] = \int_0^T (\langle p_s, \eta_s \mathcal{L}_s \eta_s^{-1} \rangle + \langle p_s, \mathcal{L}_s \log \eta_s \rangle) ds \\ &= \int_0^T (\langle \varphi_s, \mathcal{L}_s \eta_s^{-1} \rangle + \langle p_s, \mathcal{L}_s (\log \varphi_s - \log p_s) \rangle) ds \\ &\simeq \int_0^T (\langle \mathcal{L}_s^* \varphi_s, \varphi_s^{-1} p_s \rangle + \langle p_s, \mathcal{L}_s \log \varphi_s \rangle) ds \end{aligned}$$

where the notation \simeq indicates the equality up to a constant term that does not depend on $\varphi_{s|0}$. Further computation gives

$$\begin{aligned} &D_{\text{KL}}(\mathbb{P}||\mathbb{Q}) \\ &\simeq \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T (\langle \mathcal{L}_s^* \varphi_s, \varphi_s^{-1} p_{s|0}(\cdot|x_0) \rangle + \langle p_{s|0}(\cdot|x_0), \mathcal{L}_s \log \varphi_s \rangle) ds \right] \\ &= \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T (\langle \mathcal{L}_s^* \varphi_s, \eta_{s|0}^{-1}(\cdot|x_0) \rangle + \langle p_{s|0}(\cdot|x_0), \mathcal{L}_s (\log \eta_{s|0} + \log p_{s|0}) \rangle) ds \right] \\ &\simeq \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T (\langle p_{s|0}(\cdot|x_0), \eta_{s|0}(\cdot|x_0) \mathcal{L}_s \eta_{s|0}^{-1}(\cdot|x_0) \rangle + \langle p_{s|0}(\cdot|x_0), \mathcal{L}_s \log \eta_{s|0}(\cdot|x_0) \rangle) ds \right] \\ &= \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T \mathbb{E}_{x_s \sim p_{s|0}(\cdot|x_0)} \left[\eta_{s|0}(x_s|x_0) \mathcal{L}_s \eta_{s|0}^{-1}(x_s|x_0) + \mathcal{L}_s \log \eta_{s|0}(x_s|x_0) \right] ds \right], \end{aligned}$$

where the omitted terms depend only on the forward process through p_t , p_0 , and $p_{t|0}(\cdot|x_0)$. Collecting those terms yields the constant C , which does not depend on φ_t (equivalently, on the model parameters). \blacksquare

A.5 Overall Error Bounds

Proof of Theorem 3.11 We first consider one specific interval $t \in [\ell\kappa, (\ell+1)\kappa)$. Similar to the arguments in Eq. (A.9), we have, by the chain rule of KL divergence, that

$$D_{\text{KL}}(\tilde{p}_{(\ell+1)\kappa} \|\widehat{q}_{(\ell+1)\kappa}) - D_{\text{KL}}(\tilde{p}_{\ell\kappa} \|\widehat{q}_{\ell\kappa}) \leq D_{\text{KL}}(\tilde{p}_{(\ell+1)\kappa|\ell\kappa} \|\widehat{q}_{(\ell+1)\kappa|\ell\kappa} | \tilde{p}_{\ell\kappa}),$$

where the conditional KL divergence is defined as

$$D_{\text{KL}}(\tilde{p}_{(\ell+1)\kappa|\ell\kappa} \|\widehat{q}_{(\ell+1)\kappa|\ell\kappa} | \tilde{p}_{\ell\kappa}) = \mathbb{E}_{\tilde{x}_{\ell\kappa} \sim \tilde{p}_{\ell\kappa}} \left[D_{\text{KL}}(\tilde{p}_{(\ell+1)\kappa|\ell\kappa}(\cdot | \tilde{x}_{\ell\kappa}) \|\widehat{q}_{(\ell+1)\kappa|\ell\kappa}(\cdot | \tilde{x}_{\ell\kappa})) \right].$$

In the following, we will omit the subscripts for the expectation with respect to the true backward process $(\tilde{x}_t)_{t \in [0, T]}$ for brevity.

For each $\tilde{x}_{\ell\kappa} \sim \tilde{p}_{\ell\kappa}$, we consider the following decomposition of the integrand:

$$\begin{aligned} D_{\text{KL}}(\tilde{p}_{(\ell+1)\kappa|\ell\kappa}(\cdot | \tilde{x}_{\ell\kappa}) \|\widehat{q}_{(\ell+1)\kappa|\ell\kappa}(\cdot | \tilde{x}_{\ell\kappa})) &= \mathbb{E} \left[\log \frac{\tilde{p}_{(\ell+1)\kappa|\ell\kappa}(\tilde{x}_{(\ell+1)\kappa} | \tilde{x}_{\ell\kappa})}{\widehat{q}_{(\ell+1)\kappa|\ell\kappa}(\tilde{x}_{(\ell+1)\kappa} | \tilde{x}_{\ell\kappa})} \middle| \tilde{x}_{\ell\kappa} \right] \\ &= \mathbb{E} \left[\log \frac{\tilde{p}_{(\ell+1)\kappa|\ell\kappa}(\tilde{x}_{(\ell+1)\kappa} | \tilde{x}_{\ell\kappa})}{q_{(\ell+1)\kappa|\ell\kappa}(\tilde{x}_{(\ell+1)\kappa} | \tilde{x}_{\ell\kappa})} \middle| \tilde{x}_{\ell\kappa} \right] + \mathbb{E} \left[\log \frac{q_{(\ell+1)\kappa|\ell\kappa}(\tilde{x}_{(\ell+1)\kappa} | \tilde{x}_{\ell\kappa})}{\widehat{q}_{(\ell+1)\kappa|\ell\kappa}(\tilde{x}_{(\ell+1)\kappa} | \tilde{x}_{\ell\kappa})} \middle| \tilde{x}_{\ell\kappa} \right] \end{aligned}$$

where the first term is first bounded by the data-processing inequality as

$$D_{\text{KL}}(\tilde{p}_{(\ell+1)\kappa|\ell\kappa}(\cdot | \tilde{x}_{\ell\kappa}) \| q_{(\ell+1)\kappa|\ell\kappa}(\cdot | \tilde{x}_{\ell\kappa})) \leq D_{\text{KL}}(\tilde{p}_{\ell\kappa:(\ell+1)\kappa}(\cdot | \tilde{x}_{\ell\kappa}) \| q_{\ell\kappa:(\ell+1)\kappa}(\cdot | \tilde{x}_{\ell\kappa})),$$

which is further bounded by Theorem 3.6 as

$$D_{\text{KL}}(\tilde{p}_{\ell\kappa:(\ell+1)\kappa}(\cdot | \tilde{x}_{\ell\kappa}) \| q_{\ell\kappa:(\ell+1)\kappa}(\cdot | \tilde{x}_{\ell\kappa})) \leq \mathbb{E} \left[\int_{T-(\ell+1)\kappa}^{T-\ell\kappa} (\eta_t \mathcal{L}_t \eta_t^{-1} + \mathcal{L}_t \log \eta_t)(x_t) dt \middle| x_{T-\ell\kappa} \right].$$

The second term corresponding to the one-step error is bounded by Assumption 3.10 after taking expectation over $\tilde{x}_{\ell\kappa}$ as

$$\mathbb{E} \left[\mathbb{E} \left[\log \frac{q_{(\ell+1)\kappa|\ell\kappa}(\tilde{x}_{(\ell+1)\kappa} | \tilde{x}_{\ell\kappa})}{\widehat{q}_{(\ell+1)\kappa|\ell\kappa}(\tilde{x}_{(\ell+1)\kappa} | \tilde{x}_{\ell\kappa})} \middle| \tilde{x}_{\ell\kappa} \right] \right] \lesssim \kappa^{1+r},$$

and by iterating for $\ell \in [0 : L-1]$, we have

$$\begin{aligned} D_{\text{KL}}(\tilde{p}_T \|\widehat{q}_T) &= D_{\text{KL}}(\tilde{p}_{L\kappa} \|\widehat{q}_{L\kappa}) \\ &\lesssim \sum_{\ell=0}^{L-1} \left(\mathbb{E} \left[\int_{T-(\ell+1)\kappa}^{T-\ell\kappa} (\eta_t \mathcal{L}_t \eta_t^{-1} + \mathcal{L}_t \log \eta_t)(x_t) dt \right] + \kappa^{1+r} \right) + D_{\text{KL}}(\tilde{p}_0 \|\widehat{q}_0) \\ &\lesssim D_{\text{KL}}(p_T \| q_0) + \mathfrak{L}[\eta_t] + T\kappa^r, \end{aligned}$$

and the proof is complete. \blacksquare

Appendix B. Example Details

In this section, we present the proofs of the examples in the main text. Specifically, we will be verifying the assumptions that are required in Appendix A.

B.1 Diffusion Process

In this example, we choose the space $E = \mathbb{R}^d$, evidently a locally compact and separable space equipped with the Lebesgue measure. This example is the generalized case of continuous diffusion models.

Definition B.1 (Hölder Space) In \mathbb{R}^d , define the Hölder semi-norm as

$$|f|_{0,\eta} = \sup_{\mathbf{x}, \mathbf{y} \in \mathbb{R}^d, \mathbf{x} \neq \mathbf{y}} \frac{|f(\mathbf{x}) - f(\mathbf{y})|}{\|\mathbf{x} - \mathbf{y}\|^\eta},$$

the Hölder space of order k with exponent $\eta \in (0, 1]$ is defined as

$$C^{k,\eta}(\mathbb{R}^d) = \left\{ f \in C^k(\mathbb{R}^d) \mid \|f\|_{C^{k,\eta}} = \max_{|\beta| \leq k} \sup_{\mathbf{x} \in \mathbb{R}^d} |\partial_\beta f(\mathbf{x})| + \max_{|\beta|=k} |\partial_\beta f|_{0,\eta} < \infty \right\},$$

where β is a multi-index, $|\beta| = \sum_{i=1}^d \beta_i$, and $\partial_\beta = \partial^{|\beta|} / \partial x_1^{\beta_1} \cdots \partial x_d^{\beta_d}$.

Generator of Diffusion Processes. For a diffusion process $(\mathbf{x}_t)_{t \in [0, T]}$ on \mathbb{R}^d , the generator \mathcal{L}_t is given by

$$\mathcal{L}_t f = \mathbf{b}_t(\mathbf{x}) \cdot \nabla f + \frac{1}{2} \mathbf{D}_t(\mathbf{x}) : \nabla^2 f, \quad (\text{B.1})$$

where $\mathbf{b}_t(\mathbf{x}) = (b_t^i(\mathbf{x}))_{i \in [d]} \in \mathbb{R}^d$ is the drift vector, and $\mathbf{D}_t(\mathbf{x}) = (D_t^{ij}(\mathbf{x}))_{i, j \in [d]} \in \mathbb{R}^{d \times d}$ is the diffusion matrix.

The diffusion process $(\mathbf{x}_t)_{t \geq 0}$ governed by the generator \mathcal{L}_t can also be expressed in the following SDE form:

$$d\mathbf{x}_t = \mathbf{b}_t(\mathbf{x}_t) dt + \boldsymbol{\Sigma}_t(\mathbf{x}_t) d\mathbf{w}_t,$$

where the matrix $\boldsymbol{\Sigma}_t(\mathbf{x})$ satisfies $\boldsymbol{\Sigma}_t(\mathbf{x}) \boldsymbol{\Sigma}_t^\top(\mathbf{x}) = \mathbf{D}_t(\mathbf{x})$ for any $\mathbf{x} \in \mathbb{R}^d$ and $t \geq 0$, and $(\mathbf{w}_t)_{t \geq 0}$ is a standard d -dimensional Wiener process.

To proceed with further discussions, we make the following assumptions on the coefficients of the generator \mathcal{L}_t :

Assumption B.2 (Regularity of the Diffusion Process) We assume the following smoothness and regularity conditions on the coefficients of the generator \mathcal{L}_t (B.1) for any $t \in [0, T]$:

- (1) The drift vector $\mathbf{b}_t(\mathbf{x})$ satisfies that $b_t^i \in C^{1,0}(\mathbb{R}^d)$ for any $i \in [d]$ and $\mathbf{x} \in \mathbb{R}^d$;
- (2) The diffusion matrix $\mathbf{D}_t(\mathbf{x})$ is positive semidefinite and satisfies that $D_t^{ij} \in C^{2,0}(\mathbb{R}^d)$ for any $i, j \in [d]$ and $\mathbf{x} \in \mathbb{R}^d$.

Domain of the Generator. Set the function class $\mathcal{U} = C^{2,0}(\mathbb{R}^d)$, it is straightforward to see that $\mathcal{U} \subset \text{dom}(\mathcal{L}_t)$ for any $t \in [0, T]$ under Assumption B.2. Then the function class \mathcal{V} can be chosen accordingly as sufficiently rapidly decreasing functions, e.g.,

$$\mathcal{V} = \left\{ f \in C^2(\mathbb{R}^d) \mid \sup_{\mathbf{x} \in \mathbb{R}^d, |\beta| \leq 2} \left| \|\mathbf{x}\|^k \partial_\beta f(\mathbf{x}) \right| < +\infty, \forall k \in \mathbb{N} \right\},$$

such that Eq. (A.3) satisfies for any $f \in \mathcal{U}$ and $g \in \mathcal{V}$:

$$\begin{aligned}
 \langle \mathcal{L}_t f, g \rangle &= \int_{\mathbb{R}^d} g(\mathbf{x}) \mathcal{L}_t f(\mathbf{x}) d\mathbf{x} \\
 &= \int_{\mathbb{R}^d} g(\mathbf{x}) \left(\mathbf{b}_t(\mathbf{x}) \cdot \nabla f(\mathbf{x}) + \frac{1}{2} \mathbf{D}_t(\mathbf{x}) : \nabla^2 f(\mathbf{x}) \right) d\mathbf{x} \\
 &= - \int_{\mathbb{R}^d} \nabla \cdot (\mathbf{b}_t(\mathbf{x}) g(\mathbf{x})) f(\mathbf{x}) d\mathbf{x} + \int_{\mathbb{R}^d} \frac{1}{2} \nabla^2 : (\mathbf{D}_t(\mathbf{x}) g(\mathbf{x})) f(\mathbf{x}) d\mathbf{x} \\
 &= \langle f, \mathcal{L}_t^* g \rangle,
 \end{aligned}$$

where the integration by parts is justified by the choice of \mathcal{V} in the second-to-last equality, and thus

$$\mathcal{L}_t^* g = -\nabla \cdot (\mathbf{b}_t g) + \frac{1}{2} \nabla^2 : (\mathbf{D}_t g),$$

which implies that $\mathcal{V} \subset \text{dom}(\mathcal{L}_t^*)$ for any $t \in [0, T]$. It is also easy to check that for any $g \in \mathcal{V}$, $g\mathcal{U} \subset \mathcal{V}$, and thus both Assumption A.18 and Assumption A.20 are satisfied.

Backward Generator. Theorem A.22 gives the exact form of the backward generator $\tilde{\mathcal{L}}_t$ as follows:

$$\begin{aligned}
 \tilde{\mathcal{L}}_{T-t} f &= p_t^{-1} \mathcal{L}_t^*(p_t f) - p_t^{-1} f \mathcal{L}_t^* p_t \\
 &= -p_t^{-1} \nabla \cdot (\mathbf{b}_t p_t f) + \frac{1}{2} p_t^{-1} \nabla^2 : (\mathbf{D}_t p_t f) + p_t^{-1} f \nabla \cdot (\mathbf{b}_t p_t) - \frac{1}{2} p_t^{-1} f \nabla^2 : (\mathbf{D}_t p_t) \\
 &= -p_t^{-1} \nabla \cdot (\mathbf{b}_t p_t) f - p_t^{-1} \mathbf{b}_t p_t \cdot \nabla f + \frac{1}{2} p_t^{-1} \nabla \cdot (f \nabla \cdot (\mathbf{D}_t p_t)) + \frac{1}{2} p_t^{-1} (\nabla \cdot (\mathbf{D}_t p_t)) \cdot \nabla f \\
 &\quad + p_t^{-1} f \nabla \cdot (\mathbf{b}_t p_t) - \frac{1}{2} p_t^{-1} f \nabla^2 : (\mathbf{D}_t p_t) \\
 &= -\mathbf{b}_t \cdot \nabla f + \frac{1}{2} \mathbf{D}_t : \nabla^2 f + p_t^{-1} (\nabla \cdot (\mathbf{D}_t p_t)) \cdot \nabla f \\
 &= (-\mathbf{b}_t + \mathbf{D}_t \nabla \log p_t + \nabla \cdot \mathbf{D}_t) \cdot \nabla f + \frac{1}{2} \mathbf{D}_t : \nabla^2 f,
 \end{aligned}$$

where we adopted the notation

$$\nabla \cdot \mathbf{D}_t(\mathbf{x}) = \left(\sum_{j=1}^d \partial_j D_t^{ij}(\mathbf{x}) \right)_{i \in [d]},$$

corresponding to the following backward SDE:

$$d\tilde{\mathbf{x}}_t = (-\mathbf{b}_t(\tilde{\mathbf{x}}_t) + \mathbf{D}_t(\tilde{\mathbf{x}}_t) \nabla \log p_t(\tilde{\mathbf{x}}_t) + \nabla \cdot \mathbf{D}_t(\tilde{\mathbf{x}}_t)) dt + \boldsymbol{\Sigma}_t(\tilde{\mathbf{x}}_t) d\mathbf{w}_t, \quad (\text{B.2})$$

where one often defines $\mathbf{s}_t = \nabla \log p_t$ as the score function.

Similarly, Assumption A.24 reduces to the following form:

$$\mathcal{K}_{T-t} f = (-\mathbf{b}_t + \mathbf{D}_t \nabla \log \varphi_t + \nabla \cdot \mathbf{D}_t) \cdot \nabla f + \frac{1}{2} \mathbf{D}_t : \nabla^2 f,$$

and thus the Markov process $(y_t)_{t \in [0, T]}$ assumed in Assumption A.24 corresponds to another diffusion process with the following SDE:

$$d\mathbf{y}_t = (-\mathbf{b}_{T-t}(\mathbf{y}_t) + \mathbf{D}_{T-t}(\mathbf{y}_t) \hat{\mathbf{s}}_{T-t}(\mathbf{y}_t) + \nabla \cdot \mathbf{D}_{T-t}(\mathbf{y}_t)) dt + \boldsymbol{\Sigma}_{T-t}(\mathbf{y}_t) d\mathbf{w}_t,$$

where $\hat{\mathbf{s}}_t = \nabla \log \varphi_t$ is an estimate of the true score function \mathbf{s}_t .

KL Divergence and Loss Function. We also compute the KL divergence in Corollary A.30 explicitly as follows:

$$\begin{aligned}
D_{\text{KL}}(\mathbb{P}||\mathbb{Q}) &= \mathbb{E} \left[\int_0^T \eta_t(\mathbf{x}_t) \mathcal{L}_t \eta_t^{-1}(\mathbf{x}_t) + \mathcal{L}_t \log \eta_t(\mathbf{x}_t) dt \right] \\
&= \mathbb{E} \left[\int_0^T \left(\eta_t(\mathbf{x}_t) \mathbf{b}_t(\mathbf{x}_t) \cdot \nabla \eta_t^{-1}(\mathbf{x}_t) + \frac{1}{2} \eta_t(\mathbf{x}_t) \mathbf{D}_t(\mathbf{x}_t) : \nabla^2 \eta_t^{-1}(\mathbf{x}_t) \right. \right. \\
&\quad \left. \left. + \mathbf{b}_t(\mathbf{x}_t) \cdot \nabla \log \eta_t(\mathbf{x}_t) + \frac{1}{2} \mathbf{D}_t(\mathbf{x}_t) : \nabla^2 \log \eta_t(\mathbf{x}_t) \right) dt \right] \\
&= \mathbb{E} \left[\int_0^T \frac{1}{2} \mathbf{D}_t(\mathbf{x}_t) : \nabla \log \eta_t(\mathbf{x}_t) \nabla^\top \log \eta_t(\mathbf{x}_t) dt \right],
\end{aligned}$$

where the last equality is by the following identity:

$$\begin{aligned}
\eta_t \nabla^2 \eta_t^{-1} &= \eta_t \nabla (-\eta_t^{-1} \nabla \log \eta_t) \\
&= -\eta_t \left[\nabla \eta_t^{-1} \nabla^\top \log \eta_t + \eta_t^{-1} \nabla^2 \log \eta_t \right] \\
&= -\nabla^2 \log \eta_t + \nabla \log \eta_t \nabla^\top \log \eta_t.
\end{aligned}$$

Notice that

$$\nabla \log \eta_t(\mathbf{x}_t) = \nabla \log \varphi_t(\mathbf{x}_t) - \nabla \log p_t(\mathbf{x}_t) = \widehat{\mathbf{s}}_t(\mathbf{x}_t) - \nabla \log p_t(\mathbf{x}_t),$$

if we define $\mathbf{s}_t = \nabla \log p_t$ as the score function and $\widehat{\mathbf{s}}_t = \nabla \log \varphi_t$ as its estimation, then the KL divergence can be rewritten as

$$D_{\text{KL}}(\mathbb{P}||\mathbb{Q}) = \mathbb{E} \left[\int_0^T \frac{1}{2} \mathbf{D}_t(\mathbf{x}_t) : (\widehat{\mathbf{s}}_t(\mathbf{x}_t) - \nabla \log p_t(\mathbf{x}_t)) (\widehat{\mathbf{s}}_t(\mathbf{x}_t) - \nabla \log p_t(\mathbf{x}_t))^\top dt \right],$$

and the corresponding score-matching loss (Corollary A.32) is thus

$$\begin{aligned}
D_{\text{KL}}(\mathbb{P}||\mathbb{Q}) &= \mathbb{E}_{\mathbf{x}_0 \sim p_0} \left[\int_0^T \mathbb{E}_{\mathbf{x}_t \sim p_{t|0}(\cdot|\mathbf{x}_0)} \left[\right. \right. \\
&\quad \left. \left. \frac{1}{2} \mathbf{D}_t(\mathbf{x}_t) : (\widehat{\mathbf{s}}_t(\mathbf{x}_t) - \nabla \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)) (\widehat{\mathbf{s}}_t(\mathbf{x}_t) - \nabla \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0))^\top \right] dt \right] + C.
\end{aligned}$$

For diagonal diffusion matrices $\mathbf{D}_t(\mathbf{x})$, the loss function can be further simplified as

$$D_{\text{KL}}(\mathbb{P}||\mathbb{Q}) = \mathbb{E}_{\mathbf{x}_0 \sim p_0} \left[\int_0^T \mathbb{E}_{\mathbf{x}_t \sim p_{t|0}(\cdot|\mathbf{x}_0)} \left[\frac{1}{2} \text{diag } \mathbf{D}_t(\mathbf{x}_t) (\widehat{\mathbf{s}}_t(\mathbf{x}_t) - \nabla \log p_{t|0}(\mathbf{x}_t|\mathbf{x}_0))^2 \right] dt \right] + C,$$

where the square is element-wise.

B.2 Jump Process

In this example, we choose the space $E = \mathbb{X}$ to be a finite set equipped with the discrete topology and the counting measure μ . This example is the generalized case of discrete diffusion models.

Generator of Jump Processes. For a jump process $(\mathbf{x}_t)_{t \in [0, T]}$ on \mathbb{X} , the generator \mathcal{L}_t is given by

$$\mathcal{L}_t f(x) = \int_{\mathbb{X}} (f(y) - f(x)) \lambda_t(y, x) \mu(dy) = \sum_{y \in \mathbb{X}} (f(y) - f(x)) \lambda_t(y, x),$$

where $\lambda_t(\cdot, x)$ is the rate of jumps from x to y at time t .

The jump process $(\mathbf{x}_t)_{t \geq 0}$ governed by the generator \mathcal{L}_t can also be expressed in the form of a continuous-time Markov chain:

$$\frac{d}{dt} \mathbf{p}_t = \mathbf{\Lambda}_t \mathbf{p}_t, \quad \text{with} \quad \Lambda_t(y, x) = \begin{cases} \lambda_t(y, x), & \text{if } y \neq x, \\ -\sum_{y' \neq x} \lambda_t(y', x), & \text{if } y = x, \end{cases}$$

where $\mathbf{p}_t = (p_t(x))_{x \in \mathbb{X}}$ is the vector of probability masses at time t .

Domain of the Generator. The choices of the function class \mathcal{U} and \mathcal{V} are trivial in this case, as the generator \mathcal{L}_t is a linear operator on the space of real-valued functions on \mathbb{X} , and thus $\mathcal{U} = \mathcal{V} = \mathbb{R}^{|\mathbb{X}|}$ by noticing that for any $f, g \in \mathbb{R}^{|\mathbb{X}|}$, we have

$$\begin{aligned} \langle \mathcal{L}_t f, g \rangle &= \sum_{x \in \mathbb{X}} g(x) \mathcal{L}_t f(x) = \sum_{x, y \in \mathbb{X}} g(x) (f(y) - f(x)) \lambda_t(y, x) \\ &= \sum_{x, y \in \mathbb{X}} g(y) f(x) \lambda_t(x, y) - g(x) f(x) \lambda_t(x, y) \\ &= \sum_{x \in \mathbb{X}} f(x) \sum_{y \in \mathbb{X}} (g(y) \lambda_t(x, y) - g(x) \lambda_t(y, x)), \end{aligned}$$

i.e., Eq. (A.3) holds with

$$\mathcal{L}_t^* g(x) = \sum_{y \in \mathbb{X}} (g(y) \lambda_t(x, y) - g(x) \lambda_t(y, x)).$$

Backward Generator. The generator $\tilde{\mathcal{L}}_t$ of the time-reversal process $(\tilde{x}_t)_{t \in [0, T]}$ is given by Assumption A.24 as

$$\begin{aligned} \tilde{\mathcal{L}}_{T-t} f &= p_t^{-1} \mathcal{L}_t^* (p_t f) - p_t^{-1} f \mathcal{L}_t^* p_t \\ &= \frac{1}{p_t(x)} \sum_{y \in \mathbb{X}} (p_t(y) f(y) \lambda_t(x, y) - p_t(x) f(x) \lambda_t(y, x)) - \frac{f(x)}{p_t(x)} \sum_{y \in \mathbb{X}} (p_t(y) \lambda_t(x, y) - p_t(x) \lambda_t(y, x)) \\ &= \sum_{y \in \mathbb{X}} (f(y) - f(x)) \frac{p_t(y)}{p_t(x)} \lambda_t(x, y) = \sum_{y \in \mathbb{X}} (f(y) - f(x)) s_t(x, y) \lambda_t(x, y), \end{aligned}$$

where we define the score function as

$$\mathbf{s}_t(x) = (s_t(x, y))_{y \in \mathbb{X}} = \left(\frac{p_t(y)}{p_t(x)} \right)_{y \in \mathbb{X}}.$$

Thus, the backward process $(\tilde{x}_t)_{t \in [0, T]}$ also corresponds to another jump process, which can be expressed as the following continuous-time Markov chain:

$$\frac{d}{dt} \mathbf{q}_t = \bar{\mathbf{\Lambda}}_t \mathbf{q}_t, \quad \text{with} \quad \bar{\Lambda}_t(y, x) = \begin{cases} s_{T-t}(x, y) \lambda_{T-t}(x, y), & \text{if } y \neq x, \\ -\sum_{y' \neq x} s_{T-t}(x, y') \lambda_{T-t}(x, y'), & \text{if } y = x. \end{cases}$$

Assumption A.24 can be similarly reduced to

$$\mathcal{K}_{T-t}f = \sum_{y \in \mathbb{X}} (f(y) - f(x)) \frac{\varphi_t(y)}{\varphi_t(x)} \lambda_t(x, y) = \sum_{y \in \mathbb{X}} (f(y) - f(x)) \widehat{s}_t(x, y) \lambda_t(x, y),$$

where \widehat{s}_t , defined as

$$\widehat{s}_t(x) = (\widehat{s}_t(x, y))_{y \in \mathbb{X}} = \left(\frac{\varphi_t(y)}{\varphi_t(x)} \right)_{y \in \mathbb{X}},$$

is an estimate of the true score function s_t . And the estimated backward process $(y_t)_{t \in [0, T]}$ corresponds to the following continuous-time Markov chain:

$$\frac{d}{dt} \widehat{\mathbf{q}}_t = \widehat{\mathbf{\Lambda}}_t \widehat{\mathbf{q}}_t, \quad \text{with} \quad \widehat{\mathbf{\Lambda}}_t(y, x) = \begin{cases} \widehat{s}_{T-t}(x, y) \lambda_{T-t}(x, y), & \text{if } y \neq x, \\ -\sum_{y' \neq x} \widehat{s}_{T-t}(x, y') \lambda_{T-t}(x, y'), & \text{if } y = x. \end{cases}$$

KL Divergence and Loss Function. The KL divergence in Corollary A.30 can be computed explicitly as

$$\begin{aligned} D_{\text{KL}}(\mathbb{P} \parallel \mathbb{Q}) &= \mathbb{E} \left[\int_0^T \eta_t(x_t) \mathcal{L}_t \eta_t^{-1}(x_t) + \mathcal{L}_t \log \eta_t(x_t) dt \right] \\ &= \mathbb{E} \left[\int_0^T \left(\sum_{y \in \mathbb{X}} \eta_t(x_t) (\eta_t^{-1}(y) - \eta_t^{-1}(x_t)) \lambda_t(y, x_t) + \sum_{y \in \mathbb{X}} (\log \eta_t(y) - \log \eta_t(x_t)) \lambda_t(y, x_t) \right) dt \right] \\ &= \mathbb{E} \left[\int_0^T \sum_{y \in \mathbb{X}} \left(\frac{\eta_t(x_t)}{\eta_t(y)} - 1 - \log \frac{\eta_t(x_t)}{\eta_t(y)} \right) \lambda_t(y, x_t) dt \right] \\ &= \mathbb{E} \left[\int_0^T \sum_{y \in \mathbb{X}} \left(\frac{\varphi_t(x_t) p_t(y)}{\varphi_t(y) p_t(x_t)} - 1 - \log \frac{\varphi_t(x_t) p_t(y)}{\varphi_t(y) p_t(x_t)} \right) \lambda_t(y, x_t) dt \right] \\ &= \mathbb{E} \left[\int_0^T \sum_{y \in \mathbb{X}} \left(\frac{s_t(x_t, y)}{\widehat{s}_t(x_t, y)} - 1 - \log \frac{s_t(x_t, y)}{\widehat{s}_t(x_t, y)} \right) \lambda_t(y, x_t) dt \right]. \end{aligned}$$

After swapping dummy variables in the double-sum representation, this can be further transformed as follows:

$$\begin{aligned} D_{\text{KL}}(\mathbb{P} \parallel \mathbb{Q}) &= \mathbb{E} \left[\int_0^T \sum_{y \in \mathbb{X}} \left(\frac{\widehat{s}_t(x_t, y)}{s_t(x_t, y)} - 1 - \log \frac{\widehat{s}_t(x_t, y)}{s_t(x_t, y)} \right) s_t(x_t, y) \lambda_t(x_t, y) dt \right] \\ &= \int_0^T \sum_{x \in \mathbb{X}} p_t(x) \sum_{y \in \mathbb{X}} \left(\frac{\widehat{s}_t(x, y)}{s_t(x, y)} - 1 - \log \frac{\widehat{s}_t(x, y)}{s_t(x, y)} \right) s_t(x, y) \lambda_t(x, y) dt \\ &= \int_0^T \sum_{x \in \mathbb{X}} p_t(x) \sum_{y \in \mathbb{X}} \left(\frac{\widehat{s}_t(x, y)}{s_t(x, y)} - 1 - \log \frac{\widehat{s}_t(x, y)}{s_t(x, y)} \right) \frac{p_t(y)}{p_t(x)} \lambda_t(x, y) dt \\ &= \mathbb{E} \left[\int_0^T \sum_{y \in \mathbb{X}} \left(\frac{\widehat{s}_t(x_t, y)}{s_t(x_t, y)} - 1 - \log \frac{\widehat{s}_t(x_t, y)}{s_t(x_t, y)} \right) s_t(x_t, y) \lambda_t(x_t, y) dt \right]. \end{aligned}$$

By Corollary A.32, the corresponding score-matching loss is

$$\begin{aligned}
 & D_{\text{KL}}(\mathbb{P}||\mathbb{Q}) \\
 &= \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T \mathbb{E}_{x_t \sim p_{t|0}(\cdot|x_0)} \left[\sum_{y \in \mathbb{X}} \left(\frac{\varphi_t(x_t) p_{t|0}(y|x_0)}{\varphi_t(y) p_{t|0}(x_t|x_0)} - 1 - \log \frac{\varphi_t(x_t) p_{t|0}(y|x_0)}{\varphi_t(y) p_{t|0}(x_t|x_0)} \right) \lambda_t(x_t, y) \right] dt \right] + C \\
 &= \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T \sum_{x \in \mathbb{X}} p_{t|0}(x|x_0) \sum_{y \in \mathbb{X}} \left(\frac{\varphi_t(x) p_{t|0}(y|x_0)}{\varphi_t(y) p_{t|0}(x|x_0)} - \log \frac{\varphi_t(x)}{\varphi_t(y)} \right) \lambda_t(x, y) dt \right] + C \\
 &= \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T \sum_{x \in \mathbb{X}} p_{t|0}(x|x_0) \sum_{y \in \mathbb{X}} \left(\frac{\varphi_t(y) p_{t|0}(x|x_0)}{\varphi_t(x) p_{t|0}(y|x_0)} - \log \frac{\varphi_t(y)}{\varphi_t(x)} \right) \frac{p_{t|0}(y|x_0)}{p_{t|0}(x|x_0)} \lambda_t(x, y) dt \right] + C \\
 &= \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T \mathbb{E}_{x_t \sim p_{t|0}(\cdot|x_0)} \left[\sum_{y \in \mathbb{X}} \left(\frac{\varphi_t(y) p_{t|0}(x_t|x_0)}{\varphi_t(x_t) p_{t|0}(y|x_0)} - \log \frac{\varphi_t(y)}{\varphi_t(x_t)} \right) \frac{p_{t|0}(y|x_0)}{p_{t|0}(x_t|x_0)} \lambda_t(x_t, y) \right] dt \right] + C \\
 &= \mathbb{E}_{x_0 \sim p_0} \left[\int_0^T \mathbb{E}_{x_t \sim p_{t|0}(\cdot|x_0)} \left[\sum_{y \in \mathbb{X}} \left(\widehat{s}_t(x_t, y) - \frac{p_{t|0}(y|x_0)}{p_{t|0}(x_t|x_0)} \log \widehat{s}_t(x_t, y) \right) \lambda_t(x_t, y) \right] dt \right] + C.
 \end{aligned}$$

B.3 General Lévy-Type Process

We now turn to the general case of Lévy-type processes, where we take the space $E = \mathbb{R}^d$. For ease of presentation, we extend the Markov process $(\mathbf{x}_t)_{t \in [0, T]}$ to $(\mathbf{x}_t)_{t \in \mathbb{R}}$, *i.e.*, take $\mathbb{T} = \mathbb{R}$ in Definition A.9, and make the following assumption:

Assumption B.3 (Courrège) *Under Assumption A.16, the augmented process $(\tilde{\mathbf{x}}_t)_{t \in \mathbb{R}}$ is conservative, *i.e.*, $\tilde{T}_t \mathbf{1} = \mathbf{1}$, where \tilde{T}_t is the Feller semigroup defined in Definition A.6, and $C_c^\infty(\mathbb{R}^{d+1}) \subset \text{dom}(\tilde{\mathcal{L}})$, where $C_c^\infty(\mathbb{R}^{d+1})$ is the space of all compactly supported smooth functions in \mathbb{R}^{d+1} .*

Generator of General Lévy-Type Process. We first present the following theorem stating that under Assumption B.3, the most general form of the forward generator is the generator of general Lévy-Type processes.

Theorem B.4 *Under Assumptions B.3 and A.16, the right generator \mathcal{L}_s of the Markov process $(\mathbf{x}_t)_{t \in \mathbb{R}}$ at time s is of the following form:*

$$\begin{aligned}
 \mathcal{L}_s f_s(\mathbf{x}) &= \mathbf{b}_s(\mathbf{x}) \cdot \nabla f_s(\mathbf{x}) + \frac{1}{2} \mathbf{D}_s(\mathbf{x}) : \nabla^2 f_s(\mathbf{x}) \\
 &\quad + \int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (f_s(\mathbf{y}) - f_s(\mathbf{x}) - (\mathbf{y} - \mathbf{x}) \cdot \nabla f_s(\mathbf{x}) \chi_s(\mathbf{x}, \mathbf{y})) \lambda_s(d\mathbf{y}, \mathbf{x}),
 \end{aligned} \tag{B.3}$$

for any $f(\cdot) \in C_c^\infty(\mathbb{R} \times \mathbb{R}^d)$, where χ is a local unit² and the coefficients satisfy the following conditions:

-
2. We say χ is a local unit if it is a C^∞ mapping from $\mathbb{R}^d \times \mathbb{R}^d$ to $[0, 1]$ satisfying that (1) $\chi(\mathbf{x}, \mathbf{y}) = 1$ in a neighborhood of the diagonal set $\{(\mathbf{x}, \mathbf{x}) | \mathbf{x} \in \mathbb{R}^d\}$; (2) for any compact set $K \subset \mathbb{R}^d$, the mappings $\{\chi(\mathbf{x}, \cdot)\}_{\mathbf{x} \in K}$ are supported in a fixed compact set in \mathbb{R}^d .

- (1) The drift $\mathbf{b}_s(\mathbf{x}) \in \mathbb{R}^d$ for any $\mathbf{x} \in \mathbb{R}^d$;
- (2) The diffusion matrix $\mathbf{D}_s(\mathbf{x}) \in \mathbb{R}^{d \times d}$ is positive semidefinite for any $\mathbf{x} \in \mathbb{R}^d$;
- (3) $\lambda_s(\cdot, \mathbf{x})$ is a Lévy measure, i.e., a Borel measure on $\mathbb{R}^d \setminus \{\mathbf{x}\}$ and

$$\int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (1 \wedge \|\mathbf{y} - \mathbf{x}\|^2) \lambda_s(d\mathbf{y}, \mathbf{x}) < +\infty,$$

for any $\mathbf{x} \in \mathbb{R}^d$.

Proof By Theorem 3.5.5 in Applebaum (2009), which is originally derived by Courrège (1965), the augmented generator $\tilde{\mathcal{L}}$ of the augmented process $(\tilde{\mathbf{x}}_t)_{t \in \mathbb{R}}$, the domain $\text{dom}(\tilde{\mathcal{L}})$ contains $C_c^\infty(\mathbb{R}^{d+1})$ as assumed in Assumption B.3, is of the following form for any

$$\begin{aligned} \tilde{\mathcal{L}}f(\tilde{\mathbf{x}}) &= -\tilde{\gamma}(\tilde{\mathbf{x}})f(\tilde{\mathbf{x}}) + \tilde{\mathbf{b}}(\tilde{\mathbf{x}}) \cdot \tilde{\nabla}f(\tilde{\mathbf{x}}) + \frac{1}{2}\mathbf{D}(\tilde{\mathbf{x}}) : \tilde{\nabla}^2f(\tilde{\mathbf{x}}) \\ &\quad + \int_{\mathbb{R}^{d+1} \setminus \{\tilde{\mathbf{x}}\}} \left(f(\tilde{\mathbf{y}}) - f(\tilde{\mathbf{x}}) - (\tilde{\mathbf{y}} - \tilde{\mathbf{x}}) \cdot \tilde{\nabla}f(\tilde{\mathbf{x}})\chi(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \right) \lambda(d\tilde{\mathbf{y}}, \tilde{\mathbf{x}}), \end{aligned}$$

where $\tilde{\nabla} = (\partial_s, \nabla)$. The coefficients satisfy the following conditions:

- (1) The killing rate $\gamma(\tilde{\mathbf{x}}) \geq 0$ for any $\tilde{\mathbf{x}} \in \mathbb{R}^{d+1}$;
- (2) The drift $\mathbf{b}(\tilde{\mathbf{x}}) \in \mathbb{R}^{d+1}$ for any $\tilde{\mathbf{x}} \in \mathbb{R}^{d+1}$;
- (3) The diffusion matrix $\mathbf{D}(\tilde{\mathbf{x}}) \in \mathbb{R}^{(d+1) \times (d+1)}$ is positive semidefinite for any $\tilde{\mathbf{x}} \in \mathbb{R}^{d+1}$;
- (4) $\lambda(\cdot, \tilde{\mathbf{x}})$ is a Lévy measure on $\mathbb{R}^{d+1} \setminus \{\tilde{\mathbf{x}}\}$.

Now that we assume the augmented generator $\tilde{\mathcal{L}}$ is conservative, we see by the definition of the augmented generator (Definition A.11) that

$$0 = \tilde{\mathcal{L}}1 = -\tilde{\gamma} \leq 0,$$

and thus $\tilde{\gamma} = 0$ (cf. (Böttcher et al., 2013, Lemma 2.32)).

Notice that for any $f \in C_c^\infty(\mathbb{R}^{d+1})$, it also satisfies the conditions in Proposition A.12, and thus the augmented generator $\tilde{\mathcal{L}}$ only consists of the time derivative ∂_s and the right generator \mathcal{L}_s only involving operations on the function $f_s(\mathbf{x})$ at time s regarded as a function of only $\mathbf{x} \in \mathbb{R}^d$, where $f_s(\mathbf{x})$ is an alternative notation for $f(\tilde{\mathbf{x}})$ with $\tilde{\mathbf{x}} = (s, \mathbf{x})$. Therefore, the augmented generator $\tilde{\mathcal{L}}$ can be reduced to the following form:

$$\tilde{\mathcal{L}}f_s(\mathbf{x}) = \partial_s f_s(\mathbf{x}) + \mathcal{L}_s f_s(\mathbf{x}),$$

where the right generator \mathcal{L}_s satisfies

$$\begin{aligned} \mathcal{L}_s f_s(\mathbf{x}) &= \mathbf{b}_s(\mathbf{x}) \cdot \nabla f_s(\mathbf{x}) + \frac{1}{2}\mathbf{D}_s(\mathbf{x}) : \nabla^2 f_s(\mathbf{x}) \\ &\quad + \int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (f_s(\mathbf{y}) - f_s(\mathbf{x}) - (\mathbf{y} - \mathbf{x}) \cdot \nabla f_s(\mathbf{x})\chi_s(\mathbf{x}, \mathbf{y})) \lambda_s(d\mathbf{y}, \mathbf{x}), \end{aligned}$$

and it is easy to check that the required conditions of the coefficients are satisfied. The proof is thus completed. \blacksquare

We remark that the Courrège form (B.3) is closely related to the Lévy-Khintchine representation of Lévy processes, and thus, without loss of generality, we choose the local unit χ to be the indicator function $\mathbf{1}_{B(\mathbf{0},1)}$ of the unit ball $B(\mathbf{0},1)$ centered at the origin with radius 1, and give the following definition of general Lévy-type processes.

Definition B.5 (General Lévy-Type Process) *A general Lévy-type process $(\mathbf{x}_t)_{t \in \mathbb{R}}$ is a Markov process on \mathbb{R}^d with the generator \mathcal{L}_t of the following form:*

$$\begin{aligned} \mathcal{L}_t f(\mathbf{x}) = & \mathbf{b}_t(\mathbf{x}) \cdot \nabla f(\mathbf{x}) + \frac{1}{2} \mathbf{D}_t(\mathbf{x}) : \nabla^2 f(\mathbf{x}) \\ & + \int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (f(\mathbf{y}) - f(\mathbf{x}) - (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) \mathbf{1}_{B(\mathbf{x},1)}(\mathbf{y} - \mathbf{x})) \lambda_t(d\mathbf{y}, \mathbf{x}), \end{aligned} \quad (\text{B.4})$$

for any $f \in C_c^\infty(\mathbb{R}^d)$, where $\mathbf{b}_t(\mathbf{x}) \in \mathbb{R}^d$ is the drift, $\mathbf{D}_t(\mathbf{x}) \in \mathbb{R}^{d \times d}$ is the diffusion matrix, and $\lambda_t(\cdot, \mathbf{x})$ is a Lévy measure. In the following, we will call $(\mathbf{b}_t, \mathbf{D}_t, \lambda_t)$ the Lévy triplet of the general Lévy-type process.

To enable further discussion on the construction of the backward process, we summarize and strengthen the required conditions for the coefficients in the following assumption as in Assumption B.2:

Assumption B.6 (Regularity of General Lévy-Type Process) *We assume the following smoothness and regularity conditions on the Lévy triplet of the generator \mathcal{L}_t in Eq. (B.4) for any $t \in \mathbb{R}$:*

- (1) *The drift vector $\mathbf{b}_t(\mathbf{x})$ satisfies $b_t^i \in C^{1,0}(\mathbb{R}^d)$ for any $i \in [d]$ and $\mathbf{x} \in \mathbb{R}^d$;*
- (2) *The diffusion matrix $\mathbf{D}_t(\mathbf{x})$ is positive semidefinite and satisfies $D_t^{ij} \in C^{2,0}(\mathbb{R}^d)$ for any $i, j \in [d]$ and $\mathbf{x} \in \mathbb{R}^d$;*
- (3) *The Lévy measure $\lambda_t(\cdot, \mathbf{x})$ admits a density function with respect to the Lebesgue measure, which we denote with a slight abuse of notation with $\lambda_t(\mathbf{y}, \mathbf{x})$, i.e.,*

$$\lambda_t(d\mathbf{y}, \mathbf{x}) = \lambda_t(\mathbf{y}, \mathbf{x}) d\mathbf{y},$$

and the density satisfies $\lambda_t(\mathbf{x}, \mathbf{x}) = 0$, $\lambda_t(\cdot, \mathbf{x}) \in C^{1,0}(\mathbb{R}^d)$ for any $\mathbf{x} \in \mathbb{R}^d$, $\lambda_t(\mathbf{y}, \cdot) \in C^{1,0}(\mathbb{R}^d)$ for any $\mathbf{y} \in \mathbb{R}^d$, and the following integral

$$\int_{\mathbb{R}^d} \left(\frac{p_t(\mathbf{y})}{p_t(\mathbf{x})} \lambda_t(\mathbf{x}, \mathbf{y}) + \lambda_t(\mathbf{y}, \mathbf{x}) \right) (\mathbf{y} - \mathbf{x}) d\mathbf{y}$$

exists either bona fide or in the sense of Cauchy principal value, for any $\mathbf{x} \in \mathbb{R}^d$ and $t \in [0, T]$.

The following proposition gives a stochastic integral representation of the general Lévy-type process.

Proposition B.7 *Under Assumption B.6, the generator associated with the Markov process defined by following stochastic integral:*

$$\begin{aligned} \mathbf{x}_t &= \mathbf{x}_0 + \int_0^t \mathbf{b}_s(\mathbf{x}_s) ds + \int_0^t \boldsymbol{\Sigma}_s(\mathbf{x}_s) d\mathbf{w}_s \\ &+ \int_0^t \int_{\mathbb{R}^d \setminus B(\mathbf{x}_{s-}, 1)} (\mathbf{y} - \mathbf{x}_{s-}) N[\lambda](ds, d\mathbf{y}) + \int_0^t \int_{B(\mathbf{x}_{s-}, 1)} (\mathbf{y} - \mathbf{x}_{s-}) \tilde{N}[\lambda](ds, d\mathbf{y}), \end{aligned} \quad (\text{B.5})$$

coincides with the form (B.4), where the matrix $\boldsymbol{\Sigma}_s(x)$ satisfies $\boldsymbol{\Sigma}_s(x)\boldsymbol{\Sigma}_s(x)^\top = \mathbf{D}_s(x)$ for any $x \in \mathbb{R}^d$, $(\mathbf{w}_s)_{s \geq 0}$ is a d -dimensional Wiener process, $N[\lambda](ds, d\mathbf{y})$ denotes the Poisson random measure with evolving intensity $\lambda_s(\mathbf{y}, \mathbf{x}_{s-})d\mathbf{y}$, and $\tilde{N}[\lambda](ds, d\mathbf{y})$ is the compensated version of $N[\lambda](ds, d\mathbf{y})$, i.e.,

$$\tilde{N}[\lambda](ds, d\mathbf{y}) = N[\lambda](ds, d\mathbf{y}) - \lambda_s(\mathbf{y}, \mathbf{x}_{s-}) ds d\mathbf{y}.$$

Proof By taking $X(t) = \mathbf{x}_t$, $z = (\mathbf{y}, \xi)$, $\alpha(t, X(t)) = \mathbf{b}_t(\mathbf{x}_t)$, $\sigma(t, X(t)) = \boldsymbol{\Sigma}_t(\mathbf{x}_t)$, and

$$\gamma(t, z, \omega) = (\mathbf{y} - \mathbf{x}_{t-}(\omega)) \mathbf{1}_{[0, \lambda_t(\mathbf{y}, \mathbf{x}_{t-}(\omega))]}(\xi)$$

in (Øksendal and Sulem, 2019, Theorem 1.16), where $\xi \in \mathbb{R}$ is an auxiliary variable augmented to the space \mathbb{R}^d to account for varying jump intensities at different locations, following the strategy in Protter (1983), we see that the stochastic integral (B.5) corresponds to a unique strong Markov càdlàg solution and by the general form of (Øksendal and Sulem, 2019, Theorem 1.22) with the compensated Poisson random measure part, we have that

$$\begin{aligned} \mathcal{L}_t f(\mathbf{x}) &= \mathbf{b}_t(\mathbf{x}) \cdot \nabla f(\mathbf{x}) + \frac{1}{2} \mathbf{D}_t(\mathbf{x}) : \nabla^2 f(\mathbf{x}) \\ &+ \int_{\mathbb{R}^d \setminus B(\mathbf{x}, 1)} \int_{\mathbb{R}} (f(\mathbf{x} + (\mathbf{y} - \mathbf{x}) \mathbf{1}_{[0, \lambda_t(\mathbf{y}, \mathbf{x})]}(\xi)) - f(\mathbf{x})) d\mathbf{y} d\xi \\ &+ \int_{B(\mathbf{x}, 1)} \int_{\mathbb{R}} (f(\mathbf{x} + (\mathbf{y} - \mathbf{x}) \mathbf{1}_{[0, \lambda_t(\mathbf{y}, \mathbf{x})]}(\xi)) - f(\mathbf{x}) \\ &\quad - (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) \mathbf{1}_{[0, \lambda_t(\mathbf{y}, \mathbf{x})]}(\xi)) d\mathbf{y} d\xi \\ &= \mathbf{b}_t(\mathbf{x}) \cdot \nabla f(\mathbf{x}) + \frac{1}{2} \mathbf{D}_t(\mathbf{x}) : \nabla^2 f(\mathbf{x}) + \int_{\mathbb{R}^d \setminus B(\mathbf{x}, 1)} (f(\mathbf{y}) - f(\mathbf{x})) \lambda_t(\mathbf{y}, \mathbf{x}) d\mathbf{y} \\ &+ \int_{B(\mathbf{x}, 1)} (f(\mathbf{y}) - f(\mathbf{x}) - (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x})) \lambda_t(\mathbf{y}, \mathbf{x}) d\mathbf{y}, \end{aligned}$$

which is exactly the form (B.4) and the proof is complete. ■

Domain of the Generator. We set $\mathcal{U} = C^\infty(\mathbb{R}^d) \cap C^{2,0}(\mathbb{R}^d)$, and $\mathcal{V} = \mathcal{S}(\mathbb{R}^d)$, the Schwartz space on \mathbb{R}^d . Clearly, under Assumption B.6, we have that $\mathcal{U} \subset \text{dom}(\mathcal{L}_t)$ for any $t \in [0, T]$ and that for any $f \in \mathcal{U}$ and $g \in \mathcal{V}$, Eq. (A.3) holds as

$$\langle \mathcal{L}_t f, g \rangle = \int_{\mathbb{R}^d} g(\mathbf{x}) \mathcal{L}_t f(\mathbf{x}) d\mathbf{x}$$

$$\begin{aligned}
 &= \int_{\mathbb{R}^d} g(\mathbf{x}) \left(\mathbf{b}_t(\mathbf{x}) \cdot \nabla f(\mathbf{x}) + \frac{1}{2} \mathbf{D}_t(\mathbf{x}) : \nabla^2 f(\mathbf{x}) \right) d\mathbf{x} \\
 &\quad + \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} (f(\mathbf{y}) - f(\mathbf{x}) - (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) \mathbf{1}_{B(\mathbf{0},1)}(\mathbf{y} - \mathbf{x})) \lambda_t(\mathbf{y}, \mathbf{x}) g(\mathbf{x}) d\mathbf{x} d\mathbf{y} \\
 &= - \int_{\mathbb{R}^d} \nabla \cdot (\mathbf{b}_t(\mathbf{x}) g(\mathbf{x})) f(\mathbf{x}) d\mathbf{x} + \int_{\mathbb{R}^d} \frac{1}{2} \nabla^2 : (\mathbf{D}_t(\mathbf{x}) g(\mathbf{x})) f(\mathbf{x}) d\mathbf{x} \\
 &\quad + \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} (f(\mathbf{x}) \lambda_t(\mathbf{x}, \mathbf{y}) g(\mathbf{y}) - f(\mathbf{x}) \lambda_t(\mathbf{y}, \mathbf{x}) g(\mathbf{x})) d\mathbf{x} d\mathbf{y} \\
 &\quad + \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} f(\mathbf{x}) \nabla \cdot ((\mathbf{y} - \mathbf{x}) \mathbf{1}_{B(\mathbf{x},1)}(\mathbf{y} - \mathbf{x}) \lambda_t(\mathbf{y}, \mathbf{x}) g(\mathbf{x})) d\mathbf{x} d\mathbf{y} = \langle f, \mathcal{L}_t^* g \rangle,
 \end{aligned}$$

i.e., $\mathcal{V} \subset \text{dom}(\mathcal{L}_t^*)$ for any $t \in [0, T]$, with the adjoint operator \mathcal{L}_t^* given by

$$\begin{aligned}
 \mathcal{L}_t^* g &= - \nabla \cdot (\mathbf{b}_t g) + \frac{1}{2} \nabla^2 : (\mathbf{D}_t g) + \int_{\mathbb{R}^d} (\lambda_t(\mathbf{x}, \mathbf{y}) g(\mathbf{y}) - \lambda_t(\mathbf{y}, \mathbf{x}) g(\mathbf{x})) d\mathbf{y} \\
 &\quad + \int_{\mathbb{R}^d} \nabla \cdot ((\mathbf{y} - \mathbf{x}) \mathbf{1}_{B(\mathbf{x},1)}(\mathbf{y} - \mathbf{x}) \lambda_t(\mathbf{y}, \mathbf{x}) g(\mathbf{x})) d\mathbf{y}.
 \end{aligned} \tag{B.6}$$

It is also checked that for any $g \in \mathcal{V}$, $g\mathcal{U} \subset \mathcal{V}$.

Backward Generator. We plug the form of the adjoint generator \mathcal{L}_t^* (B.6) into Theorem A.22 and obtain the backward generator $\tilde{\mathcal{L}}_t$ as follows:

$$\begin{aligned}
 \tilde{\mathcal{L}}_{T-t} f &= p_t^{-1} \mathcal{L}_t^*(p_t f) - p_t^{-1} f \mathcal{L}_t p_t \\
 &= (-\mathbf{b}_t + \mathbf{D}_t \nabla \log p_t + \nabla \cdot \mathbf{D}_t) \cdot \nabla f + \frac{1}{2} \mathbf{D}_t : \nabla^2 f + \int_{\mathbb{R}^d} (f(\mathbf{y}) - f(\mathbf{x})) \frac{p_t(\mathbf{y})}{p_t(\mathbf{x})} \lambda_t(\mathbf{x}, \mathbf{y}) d\mathbf{y} \\
 &\quad + \int_{\mathbb{R}^d} (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) \mathbf{1}_{B(\mathbf{x},1)}(\mathbf{y} - \mathbf{x}) \lambda_t(\mathbf{y}, \mathbf{x}) d\mathbf{y} \\
 &= (-\mathbf{b}_t + \mathbf{D}_t \nabla \log p_t + \nabla \cdot \mathbf{D}_t) \cdot \nabla f + \frac{1}{2} \mathbf{D}_t : \nabla^2 f + \int_{\mathbb{R}^d} (f(\mathbf{y}) - f(\mathbf{x})) \frac{p_t(\mathbf{y})}{p_t(\mathbf{x})} \lambda_t(\mathbf{x}, \mathbf{y}) d\mathbf{y} \\
 &\quad + \int_{\mathbb{R}^d} (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) \mathbf{1}_{B(\mathbf{x},1)}(\mathbf{y} - \mathbf{x}) \left(\frac{p_t(\mathbf{y})}{p_t(\mathbf{x})} \lambda_t(\mathbf{x}, \mathbf{y}) + \lambda_t(\mathbf{y}, \mathbf{x}) \right) d\mathbf{y} \\
 &\quad - \int_{\mathbb{R}^d} (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) \mathbf{1}_{B(\mathbf{x},1)}(\mathbf{y} - \mathbf{x}) \frac{p_t(\mathbf{y})}{p_t(\mathbf{x})} \lambda_t(\mathbf{x}, \mathbf{y}) d\mathbf{y}
 \end{aligned}$$

which also corresponds to a general Lévy-type process with the generator of the form (B.4) with the Lévy triplet $(\tilde{\mathbf{b}}_t, \tilde{\mathbf{D}}_t, \tilde{\lambda}_t)$:

$$\begin{aligned}
 \tilde{\mathbf{b}}_{T-t}(\mathbf{x}) &= -\mathbf{b}_t(\mathbf{x}) + \mathbf{D}_t(\mathbf{x}) \nabla \log p_t(\mathbf{x}) + \nabla \cdot \mathbf{D}_t(\mathbf{x}) \\
 &\quad + \int_{B(\mathbf{x},1)} \left(\frac{p_t(\mathbf{y})}{p_t(\mathbf{x})} \lambda_t(\mathbf{x}, \mathbf{y}) + \lambda_t(\mathbf{y}, \mathbf{x}) \right) (\mathbf{y} - \mathbf{x}) d\mathbf{y}, \\
 \tilde{\mathbf{D}}_{T-t}(\mathbf{x}) &= \mathbf{D}_t(\mathbf{x}), \quad \tilde{\lambda}_{T-t}(\mathbf{y}, \mathbf{x}) = \frac{p_t(\mathbf{y})}{p_t(\mathbf{x})} \lambda_t(\mathbf{x}, \mathbf{y}).
 \end{aligned}$$

KL Divergence and Loss Function. The KL divergence in Corollary A.30 can be rewritten as

$$D_{\text{KL}}(\mathbb{P} \parallel \mathbb{Q}) = \mathbb{E} \left[\int_0^T \eta_t(\mathbf{x}_t) \mathcal{L}_t \eta_t^{-1}(\mathbf{x}_t) + \mathcal{L}_t \log \eta_t(\mathbf{x}_t) dt \right]$$

$$\begin{aligned}
&= \mathbb{E} \left[\int_0^T \frac{1}{2} \mathbf{D}_t(\mathbf{x}_t) : \nabla \log \eta_t(\mathbf{x}_t) \nabla^\top \log \eta_t(\mathbf{x}_t) dt \right. \\
&\quad \left. + \int_0^T \int_{\mathbb{R}^d} \left(\frac{\eta_t(\mathbf{x}_t)}{\eta_t(\mathbf{y})} - 1 - \log \frac{\eta_t(\mathbf{x}_t)}{\eta_t(\mathbf{y})} \right) \lambda_t(\mathbf{y}, \mathbf{x}_t) d\mathbf{y} dt \right],
\end{aligned}$$

where the first part corresponds to the diffusion part $\mathbf{b}_t(\mathbf{x}) \cdot \nabla f(\mathbf{x})$ and the second part the large jump part $\int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (f(\mathbf{y}) - f(\mathbf{x})) \lambda_t(\mathbf{y}, \mathbf{x}) d\mathbf{y}$. The small jump part $\int_{\mathbb{R}^d \setminus \{\mathbf{x}\}} (\mathbf{y} - \mathbf{x}) \cdot \nabla f(\mathbf{x}) \mathbf{1}_{B(\mathbf{x}, 1)}(\mathbf{y} - \mathbf{x}) \lambda_t(\mathbf{y}, \mathbf{x}) d\mathbf{y}$ does not contribute to the KL divergence by noticing that $\eta_t \nabla \eta_t^{-1} + \nabla \log \eta_t = 0$. Then simple algebraic manipulations lead to Eqs. (4.6) and (4.7).

Proof of Theorem 4.5 The affine linear SDE admits the explicit variation-of-constants representation

$$\mathbf{x}_T = \mathbf{M}_T \mathbf{x}_0 + \mathbf{m}_T + \mathbf{Z}_T, \quad \mathbf{m}_T := \int_0^T \mathbf{M}_{T,s} \mathbf{u}_s ds, \quad \mathbf{Z}_T := \int_0^T \mathbf{M}_{T,s} \boldsymbol{\Sigma}_s d\mathbf{w}_s,$$

where \mathbf{M}_t solves $\dot{\mathbf{M}}_t = \mathbf{A}_t \mathbf{M}_t$ with $\mathbf{M}_0 = \mathbf{I}$, and $\mathbf{M}_{T,s} := \mathbf{M}_T \mathbf{M}_s^{-1}$. For each t , \mathbf{M}_t is invertible and we denote $\sigma_{\min}(\mathbf{M}_T) > 0$ as its smallest singular value. Since $\boldsymbol{\Sigma}_t$ is bounded, $\mathbb{E} \|\mathbf{Z}_T\|^2 < \infty$ and \mathbf{m}_T is deterministic and finite.

We first show heavy-tail persistence at finite time:

$$\mathbb{E} \|\mathbf{x}_T\|^2 \geq \frac{1}{2} \mathbb{E} \|\mathbf{M}_T \mathbf{x}_0\|^2 - \mathbb{E} \|\mathbf{m}_T + \mathbf{Z}_T\|^2 \geq \frac{1}{2} \sigma_{\min}^2(\mathbf{M}_T) \mathbb{E} \|\mathbf{x}_0\|^2 - C = \infty,$$

using $\mathbb{E} \|\mathbf{x}_0\|^2 = \infty$.

Let $q_0 = \mathcal{N}(\boldsymbol{\mu}, \mathbf{Q})$ and denote by $\lambda_{\max}(\mathbf{Q})$ the largest eigenvalue. By the Donsker-Varadhan variational formula,

$$D_{\text{KL}}(p_T \| q_0) = \sup_f \left\{ \mathbb{E}_{p_T} [f(\mathbf{x})] - \log \mathbb{E}_{q_0} [e^{f(\mathbf{x})}] \right\}.$$

Choosing $f(\mathbf{x}) = \alpha \|\mathbf{x}\|^2$ with any $\alpha \in (0, \frac{1}{2\lambda_{\max}(\mathbf{Q})})$ yields

$$D_{\text{KL}}(p_T \| q_0) \geq \alpha \mathbb{E} \|\mathbf{x}_T\|^2 - \log \mathbb{E}_{q_0} [e^{\alpha \|\mathbf{x}\|^2}] = \infty,$$

which completes the proof. ■