SDDBMs: SOFT DENOISING DIFFUSION BRIDGE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion bridge models leverage Doob's h-transform to construct stochastic transports between arbitrary endpoint distributions, and have shown strong potential in image-to-image translation and restoration. However, most existing bridge models rely on hard endpoint conditioning, which forces the terminal state to match a prescribed target exactly. This hard constraint induces terminal-boundary singularities: the terminal law collapses to a Dirac measure, and the resulting drift coefficients become ill-conditioned near the endpoint. In this paper, we propose Soft Denoising Diffusion Bridge Models (SDDBMs), a generalized framework that regularizes diffusion bridges directly at the level of their terminal constraints. Instead of imposing an exact endpoint, SDDBMs prescribe a non-degenerate Gaussian terminal marginal under the transformed path measure, with a flexible terminal center and variance. Starting from this prescribed marginal, we develop a complete closed-form construction of the soft bridge, including the Gaussian terminal reweighting and soft h-function, the induced Gaussian forward marginals and $\mathbf{x}_0$-free dynamics. Theoretically, SDDBMs provide a unified probabilistic perspective that encompasses existing diffusion bridge models, including DDBMs, GOUB, and UniDB, as special cases under specific parameter choices. Extensive experiments on image restoration tasks demonstrate that SDDBMs achieve improved numerical stability and superior generation quality over existing bridge-based methods.

1 INTRODUCTION

Diffusion models have emerged as a powerful class of generative models, achieving remarkable success in image synthesis, restoration, video generation, and other domains (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2021b; Dhariwal & Nichol, 2021; Rombach et al., 2022; Podell et al., 2024). Their standard formulation gradually transports data to a simple Gaussian prior and learns the reverse denoising process. While highly effective for unconditional generation, this paradigm is less natural for tasks that require transport between two structured distributions, such as image-to-image translation and restoration. In such settings, both endpoints carry semantic and structural information, and treating one side as pure Gaussian noise may discard the intrinsic correspondence between the source and target domains.

Diffusion bridge models provide a principled alternative by constructing a stochastic process that connects two endpoint distributions. Denoising Diffusion Bridge Models (DDBMs) (Zhou et al., 2024) apply Doob's h-transform to build bridges between paired source and target domains, enabling direct distribution-to-distribution transport. GOUB (Yue et al., 2024) further extends this idea by combining Doob's h-transform with a generalized Ornstein–Uhlenbeck process, whose mean-reverting structure is particularly suitable for image restoration. These bridge-based models have demonstrated strong empirical performance on image-to-image tasks by explicitly incorporating endpoint information into the diffusion dynamics.

Despite their effectiveness, most existing Doob's h-transform bridges rely on hard endpoint conditioning: the terminal state of the process, denoted by $\mathbf{x}_T$, is forced to coincide with a prescribed target endpoint $\mathbf{x}^\star$. This hard constraint is intuitive, but it also leads to terminal-boundary singularities. For example, the bridge reverse process in DDBMs is only defined up to $T - \epsilon$, and practical sampling requires approximating the terminal state near T rather than starting exactly from the terminal boundary. Although this endpoint singularity has been observed in existing bridge formulations,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

its probabilistic origin remains insufficiently understood. In particular, it is still unclear from the viewpoint of terminal distributions why such singularities arise, whether they can be avoided, and what form of relaxation is required to remove them.

A closely related work is UniDB (Zhu et al., 2025), which revisits diffusion bridges from the perspective of stochastic optimal control (SOC). By introducing a finite terminal penalty coefficient, UniDB relaxes the infinitely strict endpoint matching of Doob's h-transform and derives a stabilized analytical control term, with hard endpoint bridges recovered as the limiting case when the terminal penalty tends to infinity. However, UniDB is primarily derived from SOC, rather than by explicitly specifying a smooth terminal reweighting under an h-transformed path measure. Consequently, while it has the potential to mitigate singularities, the probabilistic picture underlying UniDB is not very clear, and its score parameterization retains a terminal singularity requiring careful boundary treatment.

In this paper, we propose Soft Denoising Diffusion Bridge Models (SDDBMs), schematically illustrated in Figure 1, to address the terminal singularities of hard-constrained diffusion bridges at the level of terminal marginals. Our key insight is that a finite-coefficient, $\mathbf{x}_0$-free affine bridge dynamics necessarily induces a smooth, bounded, non-degenerate Gaussian terminal distribution that retains dependence on the source state $\mathbf{x}_0$. Therefore, exact source forgetting or hard endpoint collapse can only be achieved through a blow-up of the drift coefficients. To avoid this numerical singularity at its source, we replace the hard endpoint constraint with a soft terminal constraint. Specifically, instead of enforcing $\mathbf{x}_T = \mathbf{x}^\star$ exactly, we prescribe the terminal marginal under the transformed path measure as a non-degenerate Gaussian $\mathbf{x}_T \sim \mathcal{N}\left(\mathbf{x}^*, \sigma^2\mathbf{I}\right)$, where the auxiliary terminal center $\mathbf{x}^*$ may differ slightly from the desired target $\mathbf{x}^\star$ and retain a controlled dependence on $\mathbf{x}_0$. Our theoretical analysis shows that this prescribed terminal marginal can be exactly realized by Doob's h-transform through a suitable Gaussian terminal reweighting. This construction yields a closed-form soft h-function, a nonsingular analytical drift correction, and closed-form Gaussian forward marginals. It further enables a well-posed $\mathbf{x}_0$-free reformulation of the bridge dynamics. These ingredients form a complete training and sampling framework for SDDBMs, while also providing a unified probabilistic perspective on existing diffusion bridge models, including DDBMs, GOUB, and UniDB.

Our main contributions are summarized as follows:

- We prove that a finite-coefficient, $\mathbf{x}_0$-free affine bridge dynamics necessarily induces a smooth, bounded, non-degenerate Gaussian terminal distribution that still retains dependence on the source state $\mathbf{x}_0$. Consequently, within this class of dynamics, exact source forgetting or hard endpoint collapse requires the drift coefficients to blow up near the terminal time. This characterizes the numerical singularities of hard endpoint bridges from the perspective of their terminal distributions.
- To overcome the singularities of hard endpoint bridges, we introduce soft terminal constraints by prescribing a non-degenerate Gaussian terminal marginal. Starting from this prescribed marginal, we derive the corresponding Gaussian h-function in closed form, which further yields closed-form forward marginals. These results form the basis of the complete SDDBMs framework, including both the original $\mathbf{x}_0$-conditioned soft bridge and its $\mathbf{x}_0$-free reformulation.
- We show that the proposed SDDBMs framework provides a unified perspective on several existing diffusion bridge models, including DDBMs, GOUB, and UniDB. In particular, although UniDB is derived from the perspective of stochastic optimal control, our framework gives it an explicit probabilistic interpretation in terms of Gaussian terminal reweighting.
- Our model achieves state-of-the-art performance on various image-to-image tasks, including deraining, super-resolution, and inpainting.

2 PRELIMINARIES

2.1 DOOB'S h-TRANSFORM

Doob's h-transform (Särkkä & Solin, 2019; Léonard, 2011) is a classical technique for constructing conditioned or reweighted diffusion processes via a positive harmonic function. Let us consider the

reference diffusion process governed by the following stochastic differential equation:

$$\mathrm{d}\mathbf{x}_t = \mathbf{f}(\mathbf{x}_t, t)\mathrm{d}t + g_t \mathrm{d}\mathbf{w}_t, \qquad 0 \le t \le T, \tag{1}$$

where $\mathbf{x}_t \in \mathbb{R}^d$, $\mathbf{f} : \mathbb{R}^d \times [0, T] \to \mathbb{R}^d$ represents the drift coefficient, $g : [0, T] \to \mathbb{R}$ is the scalar diffusion coefficient, and $\mathbf{w}_t$ denotes a standard d-dimensional Wiener process. We assume that the reference process admits a well-defined transition density $p(\mathbf{x}_t, t \mid \mathbf{x}_s, s)$ for $0 \le s < t \le T$.

For a given positive harmonic function h, Doob's h-transform reweights the reference marginal as:

$$p_h(\mathbf{x}, t \mid \mathbf{x}_0, 0) = \frac{p(\mathbf{x}, t \mid \mathbf{x}_0, 0)h(\mathbf{x}, t)}{h(\mathbf{x}_0, 0)}. \tag{2}$$

Provided that h is sufficiently smooth, the transformed diffusion is governed by:

$$\mathrm{d}\mathbf{x}_t = \left[\mathbf{f}(\mathbf{x}_t, t) + g_t^2 \nabla_{\mathbf{x}_t} \log h(\mathbf{x}_t, t)\right] \mathrm{d}t + g_t \mathrm{d}\mathbf{w}_t, \qquad t < T. \tag{3}$$

To encode a soft terminal constraint, we introduce a nonnegative terminal weight $q : \mathbb{R}^d \to [0, \infty)$. For $0 \le t < T$, we define the h-function as its conditional expectation, with boundary condition $h(\mathbf{x}, T) = q(\mathbf{x})$. Substituting this h-function into Eq. (2) yields the modified terminal marginal density. This definition and its corresponding density are expressed jointly as:

$$h(\mathbf{x}, t) = \int_{\mathbb{R}^d} p(\mathbf{z}, T \mid \mathbf{x}, t) q(\mathbf{z})\, \mathrm{d}\mathbf{z}, \qquad p_h(\mathbf{x}, T \mid \mathbf{x}_0, 0) = \frac{p(\mathbf{x}, T \mid \mathbf{x}_0, 0) q(\mathbf{x})}{h(\mathbf{x}_0, 0)}. \tag{4}$$

Consequently, a soft terminal constraint does not arbitrarily replace the reference terminal distribution; rather, it reweights the reference terminal density in an absolutely continuous manner. Throughout, for notational simplicity, we suppress the dependence of h and p_h on the prescribed terminal constraint.

In contrast, a hard endpoint bridge imposes the singular constraint $\mathbf{x}_T = \mathbf{x}^\star$, formally corresponding to the Dirac weight $q(\mathbf{z}) = \delta(\mathbf{z} - \mathbf{x}^\star)$. For $t < T$, the hard h-function becomes $h_{\mathrm{hard}}(\mathbf{x}, t) = p(\mathbf{x}^\star, T \mid \mathbf{x}, t)$, yielding the time-$t$ marginal $p_{h_{\mathrm{hard}}}(\mathbf{x}, t \mid \mathbf{x}_0, 0) = \frac{p(\mathbf{x}, t \mid \mathbf{x}_0, 0) p(\mathbf{x}^\star, T \mid \mathbf{x}, t)}{p(\mathbf{x}^\star, T \mid \mathbf{x}_0, 0)}$, $0 < t < T$. As $t \to T$, this marginal degenerates to $\delta(\mathbf{x}_T - \mathbf{x}^\star)$.

2.2 Generalized Ornstein-Uhlenbeck Bridge

The Generalized Ornstein–Uhlenbeck (GOU) process is a flexible mean-reverting diffusion process that extends the classical Ornstein–Uhlenbeck process to time-varying coefficients (Ahmad, 1988):

$$\mathrm{d}\mathbf{x}_t = \theta_t(\boldsymbol{\mu} - \mathbf{x}_t)\, \mathrm{d}t + g_t\, \mathrm{d}\mathbf{w}_t \tag{5}$$

where $\boldsymbol{\mu} \in \mathbb{R}^d$ is the state vector, $\theta_t > 0$ and g_t are the drift and diffusion coefficients, and $\mathbf{w}_t$ is a standard d-dimensional Wiener process. Its affine drift yields the closed-form transition density $p(\mathbf{x}_t, t \mid \mathbf{x}_s, s) = \mathcal{N}\left(\mathbf{x}_t; \bar{\boldsymbol{\mu}}_{s:t}, \bar{\sigma}^2_{s:t}\mathbf{I}\right)$, $0 \le s \le t$, where

$$\bar{\boldsymbol{\mu}}_{s:t} = e^{-\bar{\theta}_{s:t}}\mathbf{x}_s + \boldsymbol{\mu}\left(1 - e^{-\bar{\theta}_{s:t}}\right), \qquad \bar{\sigma}^2_{s:t} = \int_s^t e^{-2\bar{\theta}_{u:t}} g_u^2\, du, \qquad \bar{\theta}_{s:t} := \int_s^t \theta_u\, du. \tag{6}$$

Under the common parametrization $g_t^2 = 2\lambda^2\theta_t$ with $\lambda > 0$, the transition variance reduces to $\bar{\sigma}^2_{s:t} = \lambda^2\left(1 - e^{-2\bar{\theta}_{s:t}}\right)$. Consequently, as $\bar{\theta}_{0:t} \to \infty$, the marginal converges to $\mathcal{N}(\boldsymbol{\mu}, \lambda^2\mathbf{I})$, with limiting covariance $\lambda^2\mathbf{I}$. A derivation is provided in Appendix G.

The Generalized Ornstein-Uhlenbeck Bridge (GOUB) (Yue et al., 2024) applies Doob's h-transform to the GOU process, serving as a generalization of DDBMs (Zhou et al., 2024). Despite its remarkable performance across various image-to-image tasks, its reverse-time dynamics suffer from terminal singularities. Specifically, the reverse-time SDE of GOUB is governed by the following equation:

$$\mathrm{d}\mathbf{x}_t = \left[\left(\theta_t + g_t^2 \frac{e^{-2\bar{\theta}_{t:T}}}{\bar{\sigma}^2_{t:T}}\right)(\mathbf{x}^\star - \mathbf{x}_t) - g_t^2 \nabla_{\mathbf{x}_t} \log p_h(\mathbf{x}_t, t)\right] \mathrm{d}t + g_t \mathrm{d}\bar{\mathbf{w}}_t. \tag{7}$$

However, at the terminal boundary $t = T$, Eq. (7) suffers from severe numerical singularities. First, the linear analytical drift diverges because the fraction $\frac{e^{-2\bar{\theta}_{t:T}}}{\bar{\sigma}^2_{t:T}}$ explodes as $\bar{\sigma}^2_{t:T} \to 0$. Simultaneously,

the score function $\nabla_{\mathbf{x}_t} \log p_h(\mathbf{x}_t, t)$ becomes singular, because the distribution inherently collapses to a Dirac distribution at the terminal time. Consequently, these interacting singularities prevent direct numerical sampling starting exactly from $t = T$.

UniDB (Zhu et al., 2025) revisits diffusion bridges from the perspective of stochastic optimal control and introduces a stabilized reverse-time process:

$$\mathrm{d}\mathbf{x}_t = \left[\left(\theta_t + g_t^2 \frac{e^{-2\bar{\theta}_{t:T}}}{\kappa^{-1} + \bar{\sigma}_{t:T}^2} \right) (\mathbf{x}^\star - \mathbf{x}_t) - g_t^2 \nabla_{\mathbf{x}_t} \log p_h(\mathbf{x}_t, t) \right] \mathrm{d}t + g_t \mathrm{d}\bar{\mathbf{w}}_t, \tag{8}$$

where $\kappa > 0$ is the terminal penalty coefficient. For finite κ, the additional κ^{-1} term prevents the analytical control component from diverging as $\bar{\sigma}_{t:T}^2 \to 0$, balancing endpoint accuracy against drift stability. However, UniDB is primarily derived from the perspective of SOC and does not explicitly formulate this finite-penalty mechanism as a smooth terminal reweighting under Doob's h-transform, leaving the induced terminal marginal distribution implicit. It is also worth noting that, although κ regularizes the analytical drift away from the hard-endpoint singularity, the learned score in UniDB is implemented through a noise-parameterized form $\mathbf{s}_\theta(\mathbf{x}_t, \mathbf{x}^\star, t) = -\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, \mathbf{x}^\star, t)/\bar{\sigma}'_t$, where $\mathbf{s}_\theta$ denotes the learned conditional score, $\boldsymbol{\epsilon}_\theta$ is the corresponding noise-prediction network, and $\bar{\sigma}'_t$ is the bridge noise scale used in UniDB. Since $\bar{\sigma}'_t$ vanishes at the endpoints, the endpoint behavior of this parameterization still requires careful boundary treatment in implementation.

3 Soft-Constrained Diffusion Bridges

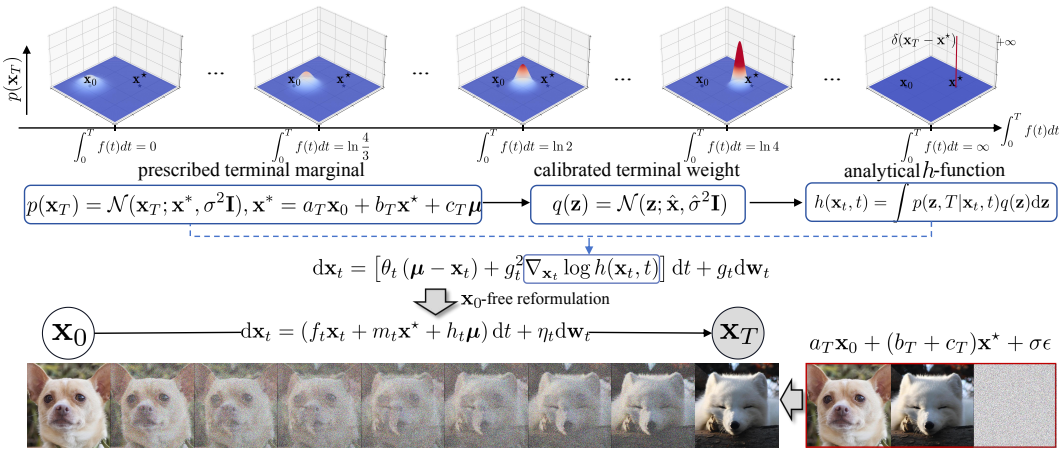

Figure 1: Illustration of hard-endpoint singularity and the SDDBMs framework. For the linear SDE $\mathrm{d}\mathbf{x}_t = f(t)(\mathbf{x}^\star - \mathbf{x}_t)\mathrm{d}t + \mathrm{d}\mathbf{w}_t$, the hard-endpoint limit $\mathbf{x}_T = \mathbf{x}^\star$ corresponds to $p(\mathbf{x}_T) \to \delta(\mathbf{x}_T - \mathbf{x}^\star)$ and requires $f(t)$ to become unbounded as $t \to T$ (see Proposition 3.1 for the general result). To avoid this degeneracy, SDDBMs prescribe a non-degenerate Gaussian terminal marginal $p(\mathbf{x}_T)$, which determines a calibrated Gaussian terminal weight $q(\mathbf{z})$ and an analytical soft h-function via terminal reweighting. The resulting h-transformed bridge admits nonsingular drift corrections and closed-form Gaussian forward marginals, and can be further rewritten into an equivalent x_0-free formulation for flexible reverse-time parameterization. The terminal-state illustration uses $\boldsymbol{\mu} = \mathbf{x}^\star$, as in our experiments.

3.1 Terminal Regularity of Finite-Coefficient Affine Bridges

We now ask what kind of terminal distributions can be generated by a finite-time bridge whose drift remains regular up to the terminal time. To make this question concrete, we consider the x_0-free affine diffusion $\mathrm{d}\mathbf{x}_t = (A_t\mathbf{x}_t + \mathbf{b}_t)\mathrm{d}t + g_t\mathrm{d}\mathbf{W}_t$, with bounded coefficients and finite nonzero diffusion energy (we relegate the complete setting description, detailed formulations of these assumptions, and their full justifications to Appendix B).

Under these assumptions, the natural question is whether a regular $\mathbf{x}_0$-free affine dynamics can produce a hard endpoint distribution or completely forget its dependence on $\mathbf{x}_0$ within a finite time

horizon. The following standard result, adapted from classical linear SDE theory (Särkkä & Solin, 2019; Oksendal, 2013), shows that neither is possible: the terminal law remains non-degenerate and retains dependence on $\mathbf{x}_0$. We formalize this result in the following proposition.

Proposition 3.1. *Consider the affine diffusion dynamics governed by the SDE* $\mathrm{d}\mathbf{x}_t = (A_t\mathbf{x}_t + \mathbf{b}_t)\mathrm{d}t + g_t\mathrm{d}\mathbf{W}_t$ *on a finite time interval* $[0,T]$*, where* $A : [0,T] \to \mathbb{R}^{d\times d}$*,* $\mathbf{b} : [0,T] \to \mathbb{R}^d$*, and* $g : [0,T] \to \mathbb{R}$ *are deterministic, time-dependent coefficients. Notably, for any time* t*, these parameters are strictly independent of the initial state* $\mathbf{x}_0$*, the current state* $\mathbf{x}_t$*, and the driving Brownian motion. Assume that*

$$\sup_{0\le t\le T}\|A_t\| < \infty, \qquad \sup_{0\le t\le T}\|\mathbf{b}_t\| < \infty, \qquad 0 < \int_0^T g_t^2\mathrm{d}t < \infty. \tag{9}$$

Let $\Phi(t,s)$ *be the fundamental matrix associated with* A_t*, defined by* $\Phi(s,s) = \mathbf{I}$*,* $\frac{\partial}{\partial t}\Phi(t,s) = A_t\Phi(t,s)$*. Then the terminal law of* $\mathbf{x}_T$ *is a non-degenerate Gaussian:* $\mathbf{x}_T \mid \mathbf{x}_0 \sim \mathcal{N}\left(\mathbf{m}_T(\mathbf{x}_0), \Sigma_T\right)$ *where* $\mathbf{m}_T(\mathbf{x}_0) = \Phi(T,0)\mathbf{x}_0 + \int_0^T \Phi(T,s)\mathbf{b}_s\,\mathrm{d}s$*,* $\Sigma_T = \int_0^T g_s^2\Phi(T,s)\Phi(T,s)^\top\,\mathrm{d}s$*. Here,* $\Phi(T,0)$ *is invertible and* Σ_T *is positive definite. Consequently,* $p(\mathbf{x},T \mid \mathbf{x}_0,0)$ *is a smooth, continuous, and bounded density. Moreover,* $\mathbf{m}_T(\mathbf{x}_0) \neq \mathbf{m}_T(\mathbf{y}_0)$ *for any* $\mathbf{x}_0 \neq \mathbf{y}_0$*.*

The proof is provided in Appendix F.1. Proposition 3.1 shows that hard endpoint bridges become singular because they require both terminal collapse to a fixed endpoint and complete forgetting of $\mathbf{x}_0$ within a finite time horizon. Further discussion is provided in Appendix C.

3.2 Regularized Diffusion Bridges with Soft Terminal Constraints

Motivated by Proposition 3.1, we regularize the bridge at the level of the terminal constraint. Rather than enforcing the endpoint through a degenerate Dirac measure, we prescribe the terminal marginal of the transformed process as a non-degenerate Gaussian:

$$p_{h_{\mathrm{soft}}}(\mathbf{x}_T, T \mid \mathbf{x}_0, 0) = \mathcal{N}\left(\mathbf{x}_T; \mathbf{x}^*, \sigma^2\mathbf{I}\right). \tag{10}$$

Here, $\mathbf{x}^*$ is a soft terminal center associated with the desired target endpoint $\mathbf{x}^\star$. In contrast to the hard bridge that enforces termination exactly at $\mathbf{x}^\star$, the proposed formulation assigns a smooth terminal density, under which endpoints closer to $\mathbf{x}^*$ have higher probability density and those farther away have lower density. Since $\mathbf{x}^*$ remains close to $\mathbf{x}^\star$, the soft terminal constraint still guides the process towards the target, while avoiding the singularity of exact endpoint conditioning.

The parameter σ controls the strength of the terminal constraint. Smaller σ yields a sharper terminal distribution and therefore a stronger preference for endpoints near $\mathbf{x}^*$, while any finite $\sigma > 0$ keeps the terminal distribution non-degenerate. As $\sigma \to 0$, the terminal marginal converges to a Dirac measure at $\mathbf{x}^*$, recovering the hard-endpoint bridge for $\mathbf{x}^* = \mathbf{x}^\star$. Thus, the proposed construction continuously relaxes the hard bridge with a finite, smooth terminal penalty.

Allowing $\mathbf{x}^* \neq \mathbf{x}^\star$ facilitates the $\mathbf{x}_0$-free reformulation by keeping the drift correction well-conditioned. We discuss the construction of $\mathbf{x}^*$ in Appendix E.2.

3.3 Terminal-Guided h-Function and Forward Process

In this section, we derive the soft h-function $h_{\mathrm{soft}}(\mathbf{x}_t, t)$ that ensures the transformed process attains the prescribed terminal marginal in Eq. (10). Using this h-function, we obtain closed-form forward marginals, which provide explicit drift corrections for the reverse-time SDE.

The following theorem derives the exact form of the terminal weighting in the h-function directly from the prescribed terminal marginals under the h-transformed measure.

Theorem 3.2. *Consider the GOU process initialized at* $\mathbf{x}_0$*, with terminal transition density* $p(\mathbf{x},T \mid \mathbf{x}_0,0) = \mathcal{N}\left(\mathbf{x}; \bar{\boldsymbol{\mu}}_{0:T}, \bar{\sigma}^2_{0:T}\mathbf{I}\right)$*, where* $\bar{\boldsymbol{\mu}}_{0:T}$ *and* $\bar{\sigma}^2_{0:T}$ *denote the conditional mean and variance of the state at time* T *given* $\mathbf{x}_0$*, respectively. Let* q *be a nonnegative probability density and define*

$$h(\mathbf{x},t) = \int p(\mathbf{z},T \mid \mathbf{x},t)q(\mathbf{z})\,\mathrm{d}\mathbf{z}. \tag{11}$$

Assume that the terminal marginal under this transformed measure is $p_h(\mathbf{x}, T \mid \mathbf{x}_0, 0) = \mathcal{N}(\mathbf{x}; \mathbf{x}^*, \sigma^2\mathbf{I})$ *for some* $0 < \sigma^2 < \bar{\sigma}^2_{0:T}$. *Then* q *is uniquely determined up to almost-everywhere equality and is given by* $q(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \hat{\mathbf{x}}, \hat{\sigma}^2\mathbf{I})$, *where*

$$\hat{\mathbf{x}} = \frac{\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2}\mathbf{x}^* - \frac{\sigma^2}{\bar{\sigma}^2_{0:T} - \sigma^2}\bar{\boldsymbol{\mu}}_{0:T}, \qquad \hat{\sigma}^2 = \frac{\sigma^2\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2}. \tag{12}$$

The proof is provided in Appendix F.2. Importantly, q is a terminal weight rather than the terminal marginal itself. Since the terminal marginal under the h-transformed path measure satisfies $p_h(\mathbf{x}, T \mid \mathbf{x}_0, 0) \propto p(\mathbf{x}, T \mid \mathbf{x}_0, 0)q(\mathbf{x})$, centering q directly at $\mathbf{x}^*$ would only yield a precision-weighted interpolation between the original terminal mean and $\mathbf{x}^*$. To realize the prescribed marginal $\mathcal{N}(\mathbf{x}; \mathbf{x}^*, \sigma^2\mathbf{I})$, we instead calibrate the terminal weight by requiring $q_{\hat{\mathbf{x}},\hat{\sigma}}(\mathbf{x}) \propto \mathcal{N}(\mathbf{x}; \mathbf{x}^*, \sigma^2\mathbf{I})/p(\mathbf{x}, T \mid \mathbf{x}_0, 0)$. For Gaussian $p(\mathbf{x}, T \mid \mathbf{x}_0, 0)$, this weight is again Gaussian provided $0 < \sigma^2 < \bar{\sigma}^2_{0:T}$, yielding Eq. (12). Thus, $\hat{\mathbf{x}}$ is a calibrated weight center rather than the terminal mean.

The following theorem provides a fully explicit closed-form expression for the soft h-function defined in Eq. (11). Furthermore, it demonstrates that this h-function satisfies the Kolmogorov Backward Equation (KBE), thereby linking the h-transform directly to the dynamics of the diffusion process.

Theorem 3.3. *Consider the transition kernel of the GOU process, given by*

$$p(\mathbf{z}, T|\mathbf{x}_t, t) = \mathcal{N}(\mathbf{z}; \bar{\boldsymbol{\mu}}_{t:T}, \bar{\sigma}^2_{t:T}\mathbf{I}), \tag{13}$$

where $\bar{\boldsymbol{\mu}}_{t:T}$ *and* $\bar{\sigma}^2_{t:T}$ *denote the conditional mean and variance of the state at time* T *given* $\mathbf{x}_t$, *respectively. Let* $q(\mathbf{z}) := \mathcal{N}(\mathbf{z}; \hat{\mathbf{x}}, \hat{\sigma}^2\mathbf{I})$ *be a Gaussian density with parameters* $\hat{\mathbf{x}} \in \mathbb{R}^d$ *and* $\hat{\sigma} > 0$. *We define the function* $h(\mathbf{x}_t, t)$ *via the integral* $h(\mathbf{x}_t, t) := \int p(\mathbf{z}, T|\mathbf{x}_t, t)q(\mathbf{z})\,\mathrm{d}\mathbf{z}$. *For any* $t \in [0, T]$, $h(\mathbf{x}_t, t)$ *admits the closed-form expression:*

$$h(\mathbf{x}_t, t) = \frac{1}{[2\pi(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})]^{d/2}} \exp\left[-\frac{\|\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}\|^2}{2(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})}\right]. \tag{14}$$

Furthermore, for any $t \in [0, T)$, $h(\mathbf{x}_t, t)$ *satisfies the Kolmogorov Backward Equation.*

The proof is provided in Appendix F.3. With Eq. (14) and Doob's h-transform, we have a new SDE:

$$\mathrm{d}\mathbf{x}_t = \left[\theta_t(\boldsymbol{\mu} - \mathbf{x}_t) + g_t^2 e^{-\bar{\theta}_{t:T}} \frac{\hat{\mathbf{x}} - \bar{\boldsymbol{\mu}}_{t:T}}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}}\right]\mathrm{d}t + g_t\mathrm{d}\mathbf{w}_t. \tag{15}$$

Compared with GOUB (Yue et al., 2024), whose drift contains the singular denominator $\bar{\sigma}^2_{t:T}$, our softened formulation introduces an additional terminal variance $\hat{\sigma}^2$. Consequently, when $\hat{\sigma}^2 > 0$, the denominator $\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}$ remains strictly positive even as $t \to T$. This prevents the analytical drift correction from exploding at the terminal boundary. Compared with UniDB (Zhu et al., 2025), our construction imposes the terminal relaxation explicitly at the level of the h-function.

Furthermore, Eq. (15) provides flexible control over the Gaussian terminal reweighting parameters $\hat{\mathbf{x}}$ and $\hat{\sigma}^2$. This flexibility yields a more stable and general diffusion bridge framework.

Recall that, for the GOU process, setting $\mathbf{x}^* = \mathbf{x}^\star$ [1] in Eq. (12) gives

$$\hat{\mathbf{x}} = -\frac{\sigma^2 e^{-\bar{\theta}_{0:T}}}{\bar{\sigma}^2_{0:T} - \sigma^2}\mathbf{x}_0 + \frac{\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2}\mathbf{x}^\star - \frac{\sigma^2\left(1 - e^{-\bar{\theta}_{0:T}}\right)}{\bar{\sigma}^2_{0:T} - \sigma^2}\boldsymbol{\mu}. \tag{16}$$

Then we can obtain an explicit parametrization of the soft h-transformed SDE in Eq. (15), which is free of $\hat{\mathbf{x}}$. See Appendix D for more details.

Given Eq. (15), the forward marginal distributions at any intermediate time $t \in [0, T]$ can be derived in closed form, yielding the following proposition:

Proposition 3.4. *Consider the stochastic process governed by Eq. (15). For* $0 < t \leq T$, *its time-*t *marginal density under the* h*-transformed path measure is*

$$p_h(\mathbf{x}_t, t|\mathbf{x}_0, 0) = \mathcal{N}\left(\mathbf{x}_t; \bar{\boldsymbol{\mu}}_{0:t} + \frac{e^{-\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t}}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}(\hat{\mathbf{x}} - \bar{\boldsymbol{\mu}}_{0:T}), \frac{\bar{\sigma}^2_{0:t}(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}\mathbf{I}\right). \tag{17}$$

The derivation of Proposition 3.4 is provided in Appendix F.4.

[1] It is important to emphasize that the restriction $\mathbf{x}^* \neq \mathbf{x}^\star$ is specific to the $\mathbf{x}_0$-free reformulation. For training or sampling formulations that condition explicitly on $\mathbf{x}_0$, one may still set $\mathbf{x}^* = \mathbf{x}^\star$.

3.4 TRACTABLE $\mathbf{x}_0$-FREE REFORMULATION

Inspired by the affine structure of $\hat{\mathbf{x}}$ parametrized by $(\mathbf{x}_0, \mathbf{x}^\star, \boldsymbol{\mu})$ in Eq. (16), we consider a more general and flexible case by setting $\hat{\mathbf{x}} = \alpha\mathbf{x}_0 + \beta\mathbf{x}^\star + \gamma\boldsymbol{\mu}$, where α, β, and γ are time-independent constants. Although certain training or sampling formulations may allow explicit conditioning on $\mathbf{x}_0$, constructing a standard reverse-time generative process where the source state is unobserved during inference typically requires a forward drift free of $\mathbf{x}_0$. By eliminating this $\mathbf{x}_0$ dependence (see Appendix E for detailed derivations), we obtain an equivalent forward process driven solely by the current state and target parameters. Following Song et al. (2021b), we have the $\mathbf{x}_0$-free reverse SDE:

$$\mathrm{d}\mathbf{x}_t = \left[f_t\mathbf{x}_t + m_t\mathbf{x}^\star + h_t\boldsymbol{\mu} - \eta_t^2\nabla_{\mathbf{x}_t}\log p_h(\mathbf{x}_t, t)\right]\mathrm{d}t + \eta_t\mathrm{d}\bar{\mathbf{w}}_t, \tag{18}$$

where

$$\phi_t = e^{-\bar{\theta}_{0:t}}\left(\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2\right) + \alpha e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}_{0:T}^2 - \bar{\sigma}_{t:T}^2\right), \qquad \varphi_t = \sigma^2\left(\bar{\sigma}_{0:T}^2 - \bar{\sigma}_{t:T}^2\right) + \bar{\sigma}_{t:T}^2\bar{\sigma}_{0:T}^2, \tag{19}$$

$$f_t = g_t^2 e^{-\bar{\theta}_{t:T}}\left(\alpha - e^{-\bar{\theta}_{0:T}}\right)/\phi_t - \theta_t, \qquad m_t = g_t^2\beta e^{-\bar{\theta}_{0:T}}/\phi_t, \tag{20}$$

$$h_t = \theta_t + g_t^2 e^{-\bar{\theta}_{t:T}}\left[e^{-\bar{\theta}_{0:T}} + (\gamma - 1)e^{-\bar{\theta}_{0:t}} + \alpha\left(e^{-\bar{\theta}_{0:t}} - 1\right)\right]\Big/\phi_t. \tag{21}$$

$$\eta_t = \frac{g_t}{\bar{\sigma}_{0:T}^2}\left\{\varphi_t\left[1 - \frac{2e^{-\bar{\theta}_{t:T}}\bar{\sigma}_{0:t}^2\left(\alpha - e^{-\bar{\theta}_{0:T}}\right)}{\phi_t}\right] + \bar{\sigma}_{0:t}^2 e^{-2\bar{\theta}_{t:T}}\left(\sigma^2 - \bar{\sigma}_{0:T}^2\right)\right\}^{1/2}. \tag{22}$$

3.5 TRAINING OBJECTIVE OF SDDBMS

Following prior diffusion bridge models (Yue et al., 2024; Zhu et al., 2025), we parameterize the reverse transition $p_\theta(\mathbf{x}_{t-1}, t-1 \mid \mathbf{x}_t, t)$ with a noise-prediction network $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, \mathbf{x}^\star, t)$. Training then matches this parameterized reverse transition with the true one-step posterior $p_h(\mathbf{x}_{t-1}, t-1 \mid \mathbf{x}_0, 0, \mathbf{x}_t, t)$. For these Gaussian reverse kernels, we train by matching their means $\boldsymbol{\mu}_{t-1}$ and $\boldsymbol{\mu}_{\theta,t-1}$:

$$\mathcal{L} = \mathbb{E}_{t,\mathbf{x}_0,\mathbf{x}_t}\left[\|\boldsymbol{\mu}_{t-1} - \boldsymbol{\mu}_{\theta,t-1}\|_1\right].$$

Here, we use the L_1 loss and omit the variance-dependent coefficient $1/(2\sigma_{\theta,t-1}^2)$ arising from the Gaussian KL objective, consistent with prior works (Yue et al., 2024; Zhu et al., 2025). Detailed derivations are provided in Appendix F.6.

4 SDDBMS UNIFY DIFFUSION BRIDGE MODELS

In this section, we show that the proposed SDDBMs framework provides a unified perspective that encompasses several existing diffusion bridge models, including DDBMs (Zhou et al., 2024), GOUB (Yue et al., 2024), and UniDB (Zhu et al., 2025).

Proposition 4.1. *The proposed SDDBMs framework recovers DDBM-VP, DDBM-VE, GOUB, and UniDB as special cases under the hyperparameter choices listed in Table 1.*

The detailed algebraic verification is provided in Appendix F.5. For DDBM-VP, DDBM-VE, and GOUB, which are based on hard constraints, the terminal marginal of the transformed process is effectively a Dirac delta distribution centered at $\mathbf{x}^\star$. Consequently, these hard-constraint methods can naturally be regarded as special cases of the soft-constraint formulation proposed in this work.

As previously discussed, although UniDB was originally derived from the perspective of stochastic optimal control, it admits a direct probabilistic interpretation under our framework. Specifically, according to the last row of Table 1, setting $\boldsymbol{\mu} = \mathbf{x}^\star$, $\hat{\mathbf{x}} = \alpha\mathbf{x}_0 + \beta\mathbf{x}^\star + \gamma\boldsymbol{\mu} = \mathbf{x}^\star$, $\hat{\sigma}^2 = \kappa^{-1}$ recovers the UniDB dynamics. Moreover, by Eq. (17), the terminal marginal of the transformed UniDB process is $p_h(\mathbf{x}_T, T \mid \mathbf{x}_0, 0) = \mathcal{N}\left(\mathbf{x}_T; \left(1 - \frac{e^{-\bar{\theta}_{0:T}}}{\kappa\bar{\sigma}_{0:T}^2+1}\right)\mathbf{x}^\star + \frac{e^{-\bar{\theta}_{0:T}}}{\kappa\bar{\sigma}_{0:T}^2+1}\mathbf{x}_0, \frac{\bar{\sigma}_{0:T}^2}{\kappa\bar{\sigma}_{0:T}^2+1}\mathbf{I}\right)$. Thus, UniDB corresponds to a soft Gaussian terminal constraint under our framework.

While UniDB is recovered as a constrained special case, SDDBMs constitute a more fundamental and flexible diffusion bridge framework rather than a mere probabilistic reinterpretation. Built upon an explicit h-transform characterization, the generalized parameters α, β, γ, and $\hat{\sigma}^2$ provide flexible control over the terminal center and variance, yielding a broader family of diffusion bridge dynamics. This flexibility inherently maintains a non-vanishing terminal noise scale, thereby avoiding the boundary training instability and specialized score parameterization required by UniDB.

Table 1: Hyper-parameter spaces for different diffusion bridge models.

$\mathcal{H}$	θ_t	$\boldsymbol{\mu}$	α	β	γ	$\hat{\sigma}^2$
DDBMs (VP)	$\frac{1}{2}g_t^2$	$\mathbf{0}$	0	1	0	0
DDBMs (VE)	0	$\mathbf{0}$	0	1	0	0
GOUB	$\frac{1}{2\lambda^2}g_t^2$	$\mathbf{x}^\star$	0	1	0	0
UniDB	$\frac{1}{2\lambda^2}g_t^2$	$\mathbf{x}^\star$	0	1	0	κ^{-1}

5 EXPERIMENTS

In this section, we evaluate our models on three popular image restoration tasks: image super-resolution, image deraining, and image inpainting. Following Zhu et al. (2025), we employ four standard evaluation metrics: Peak Signal-to-Noise Ratio (PSNR) (Fardo et al., 2016), Structural Similarity Index (SSIM) (Wang et al., 2004), Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al., 2018), and Fréchet Inception Distance (FID) (Heusel et al., 2017). We choose Bicubic (Zhu et al., 2025), MAXIM (Tu et al., 2022), MHNet (Gao et al., 2025), PromptIR (Potlapalli et al., 2023), DDRM (Kawar et al., 2022), IR-SDE (Luo et al., 2023), GOUB (Yue et al., 2024), and UniDB (Zhu et al., 2025) as our baselines for comparison. Further experimental details are provided in Appendix H.

Table 2: Quantitative comparison with relevant baselines on the DIV2K, Rain100H, and CelebA-HQ datasets. The best and second-best results are highlighted in bold and underlined, respectively.

METHOD	Image Deraining				METHOD	Image Super-Resolution				METHOD	Image Inpainting			
	PSNR↑	SSIM↑	LPIPS↓	FID↓		PSNR↑	SSIM↑	LPIPS↓	FID↓		PSNR↑	SSIM↑	LPIPS↓	FID↓
MAXIM	30.81	0.902	0.133	58.72	Bicubic	26.70	0.774	0.425	36.18	PromptIR	30.22	0.918	0.068	32.69
MHNet	31.08	0.899	0.126	57.93	DDRM	24.35	0.592	0.364	78.71	DDRM	27.16	0.899	0.089	37.02
IR-SDE	31.65	0.904	0.047	18.64	IR-SDE	25.90	0.657	0.231	45.36	IR-SDE	28.37	0.916	0.046	25.13
GOUB (SDE)	31.96	0.9028	0.046	18.14	GOUB (SDE)	26.89	0.7478	0.220	20.85	GOUB (SDE)	28.98	0.9067	0.037	4.30
GOUB (ODE)	34.56	0.9414	0.077	32.83	GOUB (ODE)	28.50	0.8070	0.328	22.14	GOUB (ODE)	31.39	0.9392	0.052	12.24
UniDB (SDE)	32.05	0.9036	0.045	17.65	UniDB (SDE)	25.46	0.6856	0.179	16.21	UniDB (SDE)	29.20	0.9077	0.036	4.08
UniDB (ODE)	34.68	0.9426	0.074	31.16	UniDB (ODE)	28.64	0.8072	0.323	22.32	UniDB (ODE)	31.67	0.9395	0.052	11.98
SDDBMs (SDE)	32.94	0.9183	**0.038**	**14.21**	SDDBMs (SDE)	27.22	0.7522	**0.136**	**13.70**	SDDBMs (SDE)	30.30	0.9201	**0.032**	**3.48**
SDDBMs (ODE)	**35.06**	**0.9464**	0.066	25.75	SDDBMs (ODE)	**28.90**	**0.8164**	0.309	19.95	SDDBMs (ODE)	**32.22**	**0.9399**	0.051	11.87

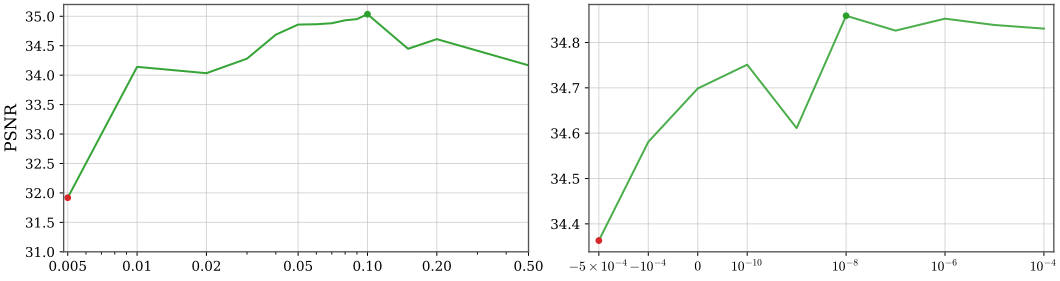

(a) Impact of σ on deraining (b) Impact of α on deraining

Figure 2: Performance sensitivity to hyperparameters σ and α on the image deraining task.

Main Results. We evaluate SDDBMs on image super-resolution, deraining, and inpainting using DIV2K (Agustsson & Timofte, 2017), Rain100H (Yang et al., 2017), and CelebA-HQ (256×256) (Karras et al., 2018), respectively. For super-resolution, the LR inputs are bicubically upsampled to match the spatial dimensions of their corresponding HR targets. For deraining, following prior works (Luo et al., 2023; Yue et al., 2024; Zhu et al., 2025), PSNR and SSIM are evaluated on the Y channel (YCbCr space). For inpainting, we adopt the 100 thin masks provided by IR-SDE (Luo et al., 2023) for testing. As reported in Table 2, SDDBMs consistently outperform the competing methods across all three tasks and evaluation metrics. The qualitative comparisons in Figure 3 further

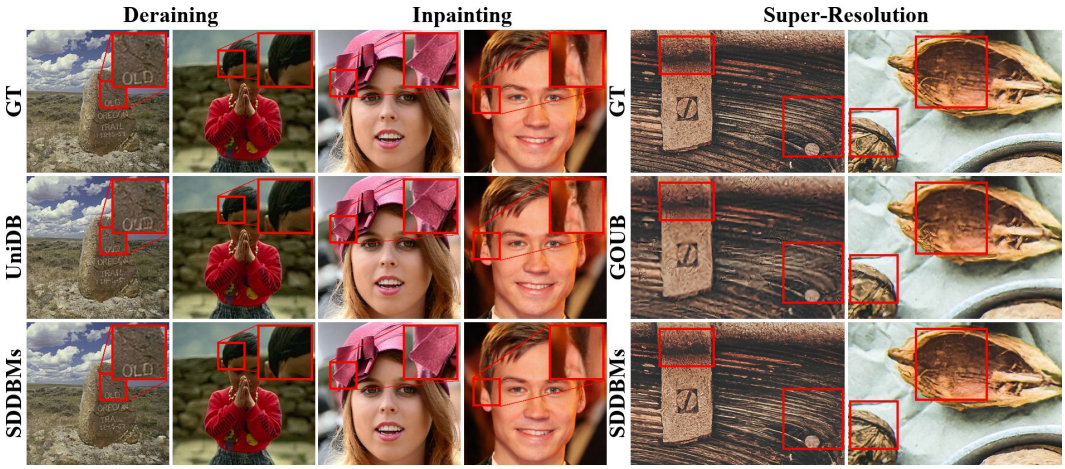

Figure 3: Qualitative comparison between SDDBMs and baseline methods on various datasets.

demonstrate that SDDBMs recover finer structural details and achieve higher visual fidelity, with additional results provided in Figures 7 and 8.

We also analyze the sensitivity of our framework to key hyperparameters: α and σ. For a detailed discussion of the motivations behind these parameter choices, together with the remaining experimental results, please refer to Appendix I.

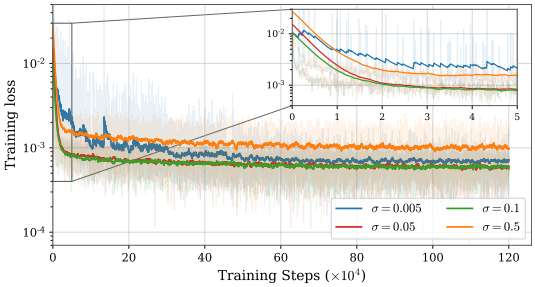

Figure 4: Training loss for different values of σ on the deraining task.

We first study the effect of the terminal variance parameter σ^2, which controls the strength of the prescribed Gaussian terminal constraint in SDDBMs. As shown in Figure 2(a), the performance first improves as σ increases but degrades when σ becomes too large. An overly small σ makes the terminal distribution highly concentrated, causing the process to approach a hard endpoint bridge and become more susceptible to terminal sampling instability, whereas an excessively large σ over-relaxes the terminal constraint and weakens target-guided transport. Figure 4 further supports this observation: both extremely small and large σ lead to higher training losses, while moderate values stabilize training. Overall, $\sigma \in [0.05, 0.11]$ provides a stable range, with the best performance typically achieved for $\sigma \in [0.08, 0.11]$. Similar trends are also observed on the other two tasks, as shown in Figure 6(a)(b).

We further investigate the effect of α. As shown in Figure 2(b), when α approaches the theoretical boundary $-e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2/\bar{\sigma}^2_{0:T}$, the model degrades markedly and becomes unstable. This performance collapse empirically validates the singularity risk proven in Lemma 2. Conversely, keeping α within an appropriate range ensures stable sampling trajectories and optimal restoration quality.

6 Conclusion

In this paper, we introduce Soft Denoising Diffusion Bridge Models (SDDBMs) to fundamentally resolve the terminal singularities inherent in hard-constrained diffusion bridges. Driven by the theoretical insight that exact endpoint collapse inevitably causes drift blow-up, our framework prescribes a non-degenerate Gaussian terminal marginal to establish soft-constrained bridge dynamics. This top-down formulation yields a fully tractable system characterized by a closed-form h-function and exact forward marginals. Building upon this probabilistic foundation, SDDBMs further unify DDBMs, GOUB, and UniDB as special cases while achieving state-of-the-art performance across diverse image-to-image translation and restoration tasks.

AI USE STATEMENT

We used ChatGPT to assist with manuscript revision and technical review, including language editing and translation, improving the organization and exposition of author-written mathematical derivations, and checking proofs, references, and author-written code for potential errors.

Beyond these editorial and review activities, generative AI was not involved in the theoretical development and proof construction or in the experimental workflow, from methodological design and implementation to data preparation, analysis, and interpretation.

All AI-assisted content was reviewed by the authors, who take full responsibility for the accuracy and integrity of the final manuscript, including its mathematical claims, implementation, and experimental results.

REFERENCES

Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 126–135, 2017.

R Ahmad. Introduction to stochastic differential equations, 1988.

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.

Julius Berner, Lorenz Richter, and Karen Ullrich. An optimal control perspective on diffusion-based generative modeling. *Trans. Mach. Learn. Res.*, 2024. URL https://openreview.net/forum?id=oYIjw37pTP.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in neural information processing systems*, 34:17695–17709, 2021.

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.

Fernando A Fardo, Victor H Conforto, Francisco C De Oliveira, and Paulo S Rodrigues. A formal evaluation of psnr as quality measurement parameter for image segmentation algorithms. *arXiv preprint arXiv:1605.07116*, 2016.

Kevin Frans, Danijar Hafner, Sergey Levine, and Pieter Abbeel. One step diffusion via shortcut models. In *International Conference on Learning Representations*, volume 2025, pp. 34668–34684, 2025.

Hu Gao, Ying Zhang, Jing Yang, and Depeng Dang. Mixed hierarchy network for image restoration. *Pattern Recognition*, 161:111313, 2025.

Zhengyang Geng, Mingyang Deng, Xingjian Bai, Zico Kolter, and Kaiming He. Mean flows for one-step generative modeling. In Danielle Belgrave, Cheng Zhang, Laura N. Montoya, Hsuan-Tien Lin, Razvan Pascanu, Piotr Koniusz, Marzyeh Ghassemi, Nancy Chen, Iván Vladimir Meza Ruíz, and Arturo Loaiza-Bonilla (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2025, NeurIPS 2025, San Diego, CA, USA, December 2-7, 2025 / Mexico City, Mexico, November 30 - December 5, 2025*, 2025. URL http://papers.nips.cc/paper_files/paper/2025/hash/6d13e085b79d454da5910e4ca82a3d9d-Abstract-Conference.html.

Jeremy Heng, Valentin De Bortoli, Arnaud Doucet, and James Thornton. Simulating diffusion bridges with score matching. *Biometrika*, 112(4):asaf048, 2025.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.

Jonathan Ho 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.

Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL https://openreview.net/forum?id=Hk99zCeAb.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35:26565–26577, 2022.

Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *Advances in neural information processing systems*, 35:23593–23606, 2022.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL http://arxiv.org/abs/1412.6980.

Christian Léonard. Stochastic derivatives and generalized h-transforms of markov processes. *arXiv preprint arXiv:1102.3172*, 2011.

Bo Li, Kaitao Xue, Bin Liu, and Yu-Kun Lai. Bbdm: Image-to-image translation with brownian bridge diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern Recognition*, pp. 1952–1961, 2023.

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, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 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, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL https://openreview.net/forum?id=XVjTT1nw5z.

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.

Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Image restoration with mean-reverting stochastic differential equations. In *International Conference on Machine Learning*, pp. 23045–23066. PMLR, 2023.

Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.

Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.

William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

Kaare Brandt Petersen and Michael Syskind Pedersen. The matrix cookbook, 2012. Version November 15, 2012.

Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. In *International Conference on Learning Representations*, volume 2024, pp. 1862–1874, 2024.

Vaishnav Potlapalli, Syed Waqas Zamir, Salman Khan, and Fahad Shahbaz Khan. Promptir: Prompting for all-in-one image restoration. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 71275–71293. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/e187897ed7780a579a0d76fd4a35d107-Paper-Conference.pdf.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.

Litu Rout, Yujia Chen, Nataniel Ruiz, Abhishek Kumar, Constantine Caramanis, Sanjay Shakkottai, and Wen-Sheng Chu. Rb-modulation: Training-free stylization using reference-based modulation. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL https://openreview.net/forum?id=bnINPG5A32.

Simo Särkkä and Arno Solin. *Applied stochastic differential equations*, volume 10. Cambridge University Press, 2019.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in neural information processing systems*, 36:62183–62223, 2023.

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

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021a. URL https://openreview.net/forum?id=St1giarCHLP.

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

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021b. URL https://openreview.net/forum?id=PxTIG12RRHS.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *International Conference on Machine Learning*, pp. 32211–32252. PMLR, 2023.

Zhengzhong Tu, Hossein Talebi, Han Zhang, Feng Yang, Peyman Milanfar, Alan Bovik, and Yinxiao Li. Maxim: Multi-axis mlp for image processing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5769–5780, 2022.

Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.

Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

Wenhan Yang, Robby T Tan, Jiashi Feng, Jiaying Liu, Zongming Guo, and Shuicheng Yan. Deep joint rain detection and removal from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1357–1366, 2017.

Conghan Yue, Zhengwei Peng, Junlong Ma, Shiyan Du, Pengxu Wei, and Dongyu Zhang. Image restoration through generalized ornstein-uhlenbeck bridge. In *International Conference on Machine Learning*, pp. 58068–58089. PMLR, 2024.

Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL https://openreview.net/forum?id=Loek7hfb46P.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.

Linqi Zhou, Aaron Lou, Samar Khanna, and Stefano Ermon. Denoising diffusion bridge models. In *International Conference on Learning Representations*, volume 2024, pp. 8160–8171, 2024.

Kaizhen Zhu, Mokai Pan, Yuexin Ma, Yanwei Fu, Jingyi Yu, Jingya Wang, and Ye Shi. Unidb: A unified diffusion bridge framework via stochastic optimal control. In *International Conference on Machine Learning*, pp. 79985–80012. PMLR, 2025.

Appendix: Soft Denoising Diffusion Bridge Models

A. Related works

Diffusion models. Since the inception of diffusion models (Sohl-Dickstein et al., 2015; Song & Ermon, 2019; Ho et al., 2020; Song et al., 2021b), various refinements have been proposed to enhance their performance. Research in this field primarily spans network design (Song et al., 2021b; Ho et al., 2020; Karras et al., 2022; Peebles & Xie, 2023), various noise schedules (Nichol & Dhariwal, 2021; Karras et al., 2022), and sampling acceleration. The latter includes both training-free numerical solvers (Lu et al., 2022; Song et al., 2021a; Zhang & Chen, 2023) and training-based methods (Song et al., 2023; Frans et al., 2025; Geng et al., 2025). Despite these successes, standard diffusion models typically rely on Gaussian priors, which limits their flexibility in tasks requiring specific endpoint-to-endpoint mapping.

Conditional generation. Approaches for conditional synthesis generally follow two primary trajectories. The first category focuses on modeling the conditional probability $p(\mathbf{x}|\mathbf{y})$, a paradigm widely adopted in scenarios such as text-to-image generation (Dhariwal & Nichol, 2021; Ho & Salimans, 2022; Rombach et al., 2022). The second category directly models a transport between two distinct distributions. For instance, inspired by the continuity equation, flow matching (Lipman et al., 2023; Albergo et al., 2025; Liu et al., 2023) learns a transport map between two distributions. Schrödinger bridges (Shi et al., 2023; De Bortoli et al., 2021) seek an optimal stochastic process π^* for inter-distribution transport under KL divergence minimization. However, the computational cost of the iterative procedures limits their practical applicability. Alternatively, diffusion bridge models establish deterministic endpoint processes by combining linear SDEs with Doob's h-transform. While successful in image-to-image translation, these models (Yue et al., 2024; Zhou et al., 2024; Li et al., 2023) encounter numerical singularities at the terminal boundary which prevent direct sampling from the terminal time—a critical limitation that our work aims to resolve.

Diffusion models with Stochastic Optimal Control. In recent years, many studies have sought to integrate SOC with diffusion models. DIS (Berner et al., 2024) established a connection between stochastic optimal control and diffusion models, potentially enabling the transfer of methods from optimal control theory to generative modeling. Rb-modulation (Rout et al., 2025) built on concepts from stochastic optimal control to modulate the drift field of reverse diffusion dynamics for training-free style transfer. UniDB (Zhu et al., 2025) offered a new perspective on diffusion bridges based on SOC which is also a special case of our SDDBMs.

B. Complete Setting and Detailed Justifications for Affine Diffusion Dynamics

We now ask what kind of terminal distribution can be generated by a finite-time bridge dynamics if its drift remains numerically regular up to the terminal time. To make this question concrete, we focus on affine diffusion dynamics of the form:

$$\mathrm{d}\mathbf{x}_t = (A_t\mathbf{x}_t + \mathbf{b}_t)\mathrm{d}t + g_t\mathrm{d}\mathbf{W}_t.$$

This class is particularly relevant for diffusion bridge models, because many widely used bridge constructions have an affine analytic bridge drift, or an affine h-drift correction, at the level of their closed-form bridge term. Examples include Brownian bridges, GOU/GOUB-type bridges, UniDB-GOU-type finite-penalty bridges, and many DDBMs bridge kernels. Although the full reverse-time sampler may additionally contain a learned nonlinear score network, the explicit terminal bridge correction often takes the affine form $\mathbf{B}_h(\mathbf{x}, t) = A_t\mathbf{x} + \mathbf{b}_t$. Thus, affine dynamics provide a useful setting in which the source of terminal instability can be isolated through the time-dependent coefficients A_t and $\mathbf{b}_t$.

We further consider the common setting where A_t, $\mathbf{b}_t$, and g_t are independent of the source state $\mathbf{x}_0$. Such an x_0-free forward formulation is often desirable when deriving the corresponding reverse-time dynamics, since the forward process can then be specified without explicit access to the source sample. We also assume that $0 < \int_0^T g_t^2 \,\mathrm{d}t < \infty$. This condition ensures that the dynamics injects a nonzero but finite amount of stochastic noise: the terminal distribution is therefore non-deterministic, while

the total diffusion energy remains finite. Finally, to exclude drift-level singularity, we require the affine coefficients to remain finite up to the terminal boundary:

$$\sup_{0\le t\le T}\|A_t\|<\infty, \qquad \sup_{0\le t\le T}\|\mathbf{b}_t\|<\infty.$$

This condition excludes coefficient-level singularities anywhere on the finite interval $[0,T]$. This requirement is directly tied to neural parameterization and numerical stability. When these coefficients are to be represented or approximated by a neural network, a finite-parameter network with finite weights can only produce finite coefficient values on the compact time interval $[0,T]$. If the target coefficients become unbounded at any time, the network cannot represent them uniformly without leaving the finite-valued regime, and the resulting dynamics may suffer from unstable training, exploding gradients, or amplified sampling errors. The terminal-time blow-up in hard endpoint bridges provides a canonical example of this coefficient-level singularity.

C. Singularities of Exact Endpoint Conditioning

As discussed in Section 2.1, conventional diffusion bridge models (Yue et al., 2024) construct a hard endpoint bridge by choosing the h-function as the transition density to a prescribed terminal point $\mathbf{x}^\star$, given by $h_{\mathrm{hard}}(\mathbf{x}_t,t)=\int\delta(\mathbf{z}-\mathbf{x}^\star)p(\mathbf{z},T\mid\mathbf{x}_t,t)\,\mathrm{d}\mathbf{z}$. This enforces exact terminal conditioning, requiring the process to hit $\mathbf{x}^\star$ at time T (Särkkä & Solin, 2019; Heng et al., 2025), which causes the terminal law to degenerate to a Dirac measure $p_{h_{\mathrm{hard}}}(\mathbf{x}_T,T\mid\mathbf{x}_0,0)=\delta(\mathbf{x}_T-\mathbf{x}^\star)$.

Proposition 3.1 shows that a finite-coefficient affine bridge dynamics yields a non-degenerate Gaussian terminal law whose density is smooth, continuous, bounded, and still depends on the source state $\mathbf{x}_0$. Exact endpoint conditioning violates both properties, leading to two closely related singularities.

Singularity of exact noise cancellation. A defining characteristic of diffusion processes is the continuous injection of stochastic noise. Under the reference dynamics, this perpetual randomness naturally yields a terminal distribution with strictly positive variance. In stark contrast, exact endpoint conditioning requires the process to reach a precise, deterministic state $\mathbf{x}^\star$, forcing the terminal variance to collapse exactly to zero. Because the diffusion term continuously injects noise up to the final moment T, the only way to perfectly counteract this inherent randomness is to apply an overwhelmingly strong deterministic correction. Consequently, to completely suppress the local stochasticity and force a Dirac terminal law, the effective drift must inevitably blow up to infinity as $t\to T$.

Finite-time source-forgetting singularity. The hard constraint also imposes exact source forgetting: regardless of $\mathbf{x}_0$, all admissible paths are forced to terminate at the same endpoint $\mathbf{x}^\star$. In contrast, Proposition 3.1 shows that, within finite-time $\mathbf{x}_0$-free affine bridge dynamics, finite drift coefficients necessarily preserve source dependence through the terminal mean. Indeed, the only source-dependent term in the terminal mean is $\Phi(T,0)\mathbf{x}_0$. Exact source forgetting would require the source-propagation map $\Phi(T,0)$ to degenerate to the zero map, whereas bounded affine coefficients make $\Phi(T,0)$ invertible. Thus, within this affine $\mathbf{x}_0$-free setting, hard endpoint collapse cannot occur in the finite-coefficient regime and can only arise through a terminal-time blow-up of the drift coefficients.

These singularities manifest as terminal drift blow-up. As $t\to T$, both the conditional variance $\bar\sigma^2_{t:T}$ of the reference GOU process and the time-t variance of the corresponding hard endpoint bridge vanish, with the bridge marginal converging to a Dirac measure. This produces sharp or unbounded score targets, amplifies discretization errors, and places a severe burden on finite-capacity neural networks and finite-step samplers.

D. Explicit Parametrization

By Eqs. (15) and (16), we have

$$\begin{aligned}\mathrm{d}\mathbf{x}_t=&\Bigg\{-\left(\theta_t+g_t^2\frac{e^{-2\bar\theta_{t:T}}}{\hat\sigma^2+\bar\sigma^2_{t:T}}\right)\mathbf{x}_t+\left[\theta_t+g_t^2\frac{e^{-\bar\theta_{t:T}}}{\hat\sigma^2+\bar\sigma^2_{t:T}}\left(e^{-\bar\theta_{t:T}}-\frac{\sigma^2\left(1-e^{-\bar\theta_{0:T}}\right)}{\bar\sigma^2_{0:T}-\sigma^2}-1\right)\right]\boldsymbol{\mu}\\&+g_t^2\frac{e^{-\bar\theta_{t:T}}}{\hat\sigma^2+\bar\sigma^2_{t:T}}\left(-\frac{\sigma^2e^{-\bar\theta_{0:T}}}{\bar\sigma^2_{0:T}-\sigma^2}\mathbf{x}_0+\frac{\bar\sigma^2_{0:T}}{\bar\sigma^2_{0:T}-\sigma^2}\mathbf{x}^\star\right)\Bigg\}\mathrm{d}t+g_t\mathrm{d}\mathbf{w}_t.\end{aligned}\tag{23}$$

Eq. (23) is an explicit parametrization of the soft h-transformed SDE in Eq. (15). It gives a drift that is affine in $\mathbf{x}_t$ and depends explicitly only on the source state $\mathbf{x}_0$, the target state $\mathbf{x}^\star$, and the prior mean $\boldsymbol{\mu}$. Its parameters are chosen so that the terminal marginal under the h-transformed path measure is exactly $p_h(\mathbf{x}_T, T \mid \mathbf{x}_0, 0) = \mathcal{N}\left(\mathbf{x}_T; \mathbf{x}^\star, \sigma^2\mathbf{I}\right)$.

E. $\mathbf{x}_0$-FREE REFORMULATION

Recall that $\hat{\mathbf{x}} = \alpha\mathbf{x}_0 + \beta\mathbf{x}^\star + \gamma\boldsymbol{\mu}$, where α, β, and γ are constants independent of t. By Eq. (6), we have $\hat{\mathbf{x}} - \bar{\boldsymbol{\mu}}_{t:T} = \alpha\mathbf{x}_0 + \beta\mathbf{x}^\star - e^{-\bar{\theta}_{t:T}}\mathbf{x}_t + \left(e^{-\bar{\theta}_{t:T}} + \gamma - 1\right)\boldsymbol{\mu}$. Consequently, Eq. (15) becomes

$$\begin{aligned}\mathrm{d}\mathbf{x}_t = &\Bigg\{ -\left(\theta_t + g_t^2 \frac{e^{-2\bar{\theta}_{t:T}}}{\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2}\right)\mathbf{x}_t + \left[\theta_t + g_t^2 \frac{e^{-\bar{\theta}_{t:T}}}{\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2}\left(e^{-\bar{\theta}_{t:T}} + \gamma - 1\right)\right]\boldsymbol{\mu} \\ &+ g_t^2 \frac{e^{-\bar{\theta}_{t:T}}}{\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2}\left(\alpha\mathbf{x}_0 + \beta\mathbf{x}^\star\right)\Bigg\}\mathrm{d}t + g_t\mathrm{d}\mathbf{w}_t.\end{aligned} \tag{24}$$

The purpose of this section is to derive an $\mathbf{x}_0$-free reformulation of the soft h-transformed dynamics in Eq. (24). Although Eq. (24) is well-defined and nonsingular, its drift explicitly depends on the initial state $\mathbf{x}_0$. For deriving the corresponding reverse-time dynamics, it is often desirable to express the forward process only in terms of the current state $\mathbf{x}_t$, the target state $\mathbf{x}^\star$, and the prior mean $\boldsymbol{\mu}$. To this end, we construct an equivalent linear SDE whose drift no longer contains $\mathbf{x}_0$ explicitly.

We perform this reformulation in two steps. First, we analyze the deterministic flow induced by the drift of Eq. (24). The closed-form flow allows us to express $\mathbf{x}_0$ as a function of $(\mathbf{x}_t, \mathbf{x}^\star, \boldsymbol{\mu})$, which determines the drift coefficients of the $\mathbf{x}_0$-free SDE. Second, since modifying the drift also changes the variance evolution of the process, we calibrate the diffusion coefficient η_t so that the time marginals of the reformulated SDE remain consistent with the h-transformed marginal distribution in Proposition 3.4.

As shown in Proposition 3.1, any finite-time affine bridge with non-singular, initial-state-independent coefficients inherently retains a dependence on the source state $\mathbf{x}_0$. We therefore introduce the ρ-relaxed terminal matching $\mathbf{x}^* = \mathbf{x}^\star + \rho(\mathbf{x}_0 - \mathbf{x}^\star)$, $\rho > 0$. Instead of forcing the terminal Gaussian center to be exactly $\mathbf{x}^\star$ and fully independent of $\mathbf{x}_0$, this relaxation allows a controlled source dependence. For small ρ, $\mathbf{x}^*$ remains close to $\mathbf{x}^\star$, so the bridge is still guided toward the target while avoiding the finite-time source-forgetting singularity caused by exact terminal matching.

E.1. $\mathbf{x}_0$-FREE DRIFT REFORMULATION

We first consider the deterministic dynamics obtained by removing the stochastic term from Eq. (24). The resulting trajectory admits the following closed-form solution.

Proposition 1. *The deterministic trajectory obtained by removing the Brownian noise term from Eq. (24) admits the following closed-form solution:*

$$\begin{aligned}\mathbf{x}_t = &\frac{e^{-\bar{\theta}_{0:t}}\left(\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2\right) + \alpha e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}_{0:T}^2 - \bar{\sigma}_{t:T}^2\right)}{\hat{\sigma}^2 + \bar{\sigma}_{0:T}^2}\mathbf{x}_0 + \frac{\beta e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}_{0:T}^2 - \bar{\sigma}_{t:T}^2\right)}{\hat{\sigma}^2 + \bar{\sigma}_{0:T}^2}\mathbf{x}^\star \\ &+ \left[1 - \frac{e^{-\bar{\theta}_{0:t}}\left(\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2\right) + (1-\gamma)e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}_{0:T}^2 - \bar{\sigma}_{t:T}^2\right)}{\hat{\sigma}^2 + \bar{\sigma}_{0:T}^2}\right]\boldsymbol{\mu}.\end{aligned} \tag{25}$$

Proof. Consider the stochastic process $\mathbf{x}_t$ governed by Eq. (24). Define $\mathbf{m}_t := \mathbb{E}[\mathbf{x}_t]$. Since Eq. (24) is linear in $\mathbf{x}_t$ and has an additive diffusion coefficient independent of $\mathbf{x}_t$, the deterministic trajectory obtained by removing the Brownian noise term from Eq. (24) is precisely $\mathbf{m}_t$. Moreover, Eq. (17) implies

$$\mathbf{x}_t \sim \mathcal{N}\left(\bar{\boldsymbol{\mu}}_{0:t} + \frac{e^{-\bar{\theta}_{t:T}}\bar{\sigma}_{0:t}^2}{\hat{\sigma}^2 + \bar{\sigma}_{0:T}^2}(\hat{\mathbf{x}} - \bar{\boldsymbol{\mu}}_{0:T}), \frac{\bar{\sigma}_{0:t}^2(\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2)}{\hat{\sigma}^2 + \bar{\sigma}_{0:T}^2}\mathbf{I}\right). \tag{26}$$

Hence,

$$\mathbf{m}_t = \bar{\boldsymbol{\mu}}_{0:t} + \frac{e^{-\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t}}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}} \left(\hat{\mathbf{x}} - \bar{\boldsymbol{\mu}}_{0:T}\right). \tag{27}$$

Recall that

$$\hat{\mathbf{x}} = \alpha \mathbf{x}_0 + \beta \mathbf{x}^\star + \gamma \boldsymbol{\mu}. \tag{28}$$

Moreover, for the GOU process,

$$\bar{\boldsymbol{\mu}}_{0:T} = e^{-\bar{\theta}_{0:T}}\mathbf{x}_0 + \left(1 - e^{-\bar{\theta}_{0:T}}\right)\boldsymbol{\mu}, \qquad \bar{\sigma}^2_{0:T} = e^{-2\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t} + \bar{\sigma}^2_{t:T}. \tag{29}$$

For compactness, define

$$K_t := \frac{e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right)}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}. \tag{30}$$

Substituting Eqs. (28) and (29) into Eq. (27) gives

$$\begin{aligned}\mathbf{m}_t =& e^{-\bar{\theta}_{0:t}}\mathbf{x}_0 + \left(1 - e^{-\bar{\theta}_{0:t}}\right)\boldsymbol{\mu} \\ &+ K_t\Big[\alpha\mathbf{x}_0 + \beta\mathbf{x}^\star + \gamma\boldsymbol{\mu} - e^{-\bar{\theta}_{0:T}}\mathbf{x}_0 - \left(1 - e^{-\bar{\theta}_{0:T}}\right)\boldsymbol{\mu}\Big].\end{aligned} \tag{31}$$

Hence,

$$\begin{aligned}\mathbf{m}_t =& \left[e^{-\bar{\theta}_{0:t}} + K_t\left(\alpha - e^{-\bar{\theta}_{0:T}}\right)\right]\mathbf{x}_0 + \beta K_t \mathbf{x}^\star \\ &+ \left[1 - e^{-\bar{\theta}_{0:t}} + K_t\left(\gamma - 1 + e^{-\bar{\theta}_{0:T}}\right)\right]\boldsymbol{\mu}.\end{aligned} \tag{32}$$

Using $e^{-\bar{\theta}_{0:T}} = e^{-\bar{\theta}_{0:t}}e^{-\bar{\theta}_{t:T}}$, we have

$$K_t e^{-\bar{\theta}_{0:T}} = e^{-\bar{\theta}_{0:t}}\frac{\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}. \tag{33}$$

Therefore,

$$e^{-\bar{\theta}_{0:t}} - K_t e^{-\bar{\theta}_{0:T}} = e^{-\bar{\theta}_{0:t}}\frac{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}. \tag{34}$$

Substituting Eq. (34) and the definition of K_t into Eq. (32) yields

$$\begin{aligned}\mathbf{m}_t =& \frac{e^{-\bar{\theta}_{0:t}}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right) + \alpha e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right)}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}\mathbf{x}_0 + \frac{\beta e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right)}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}\mathbf{x}^\star \\ &+ \left[1 - \frac{e^{-\bar{\theta}_{0:t}}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right) + (1-\gamma)e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right)}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}\right]\boldsymbol{\mu}.\end{aligned} \tag{35}$$

The proposition is proved. □

Rearranging Eq. (25), we obtain

$$\begin{aligned}\mathbf{x}_0 =& \frac{\left(\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}\right)\mathbf{x}_t - \beta e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right)\mathbf{x}^\star}{\phi_t} \\ &- \frac{\left[\left(\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}\right) - e^{-\bar{\theta}_{0:t}}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right) - (1-\gamma)e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right)\right]\boldsymbol{\mu}}{\phi_t},\end{aligned} \tag{36}$$

where ϕ_t is defined in Eq. (19).

Substituting $\mathbf{x}_0$ into the drift of Eq. (24) gives the reformulated forward dynamics

$$\mathrm{d}\mathbf{x}_t = \left(f_t\mathbf{x}_t + m_t\mathbf{x}^\star + h_t\boldsymbol{\mu}\right)\mathrm{d}t + \eta_t \mathrm{d}\mathbf{w}_t, \tag{37}$$

where η_t is an undetermined diffusion coefficient, and f_t, m_t and h_t are defined in Eqs. (20) and (21), respectively.

E.2. Well-Posedness via ρ-Relaxed Terminal Matching

Note that Eq. (37) becomes ill-defined whenever $\phi_t = 0$, since the drift coefficients contain $1/\phi_t$. Therefore, to ensure that the reformulated dynamics remain well-defined and free of such singularities, we require $\phi_t \neq 0$ over the time interval of interest. The following lemma gives an equivalent characterization of this condition. Recall that $\hat{\mathbf{x}} = \alpha\mathbf{x}_0 + \beta\mathbf{x}^\star + \gamma\boldsymbol{\mu}$ and $\mathbf{x}^* = \mathbf{x}^\star + \rho(\mathbf{x}_0 - \mathbf{x}^\star)$.

Lemma 2. *The following statements are equivalent: (i) $\phi_t \neq 0$ for all $t \in [0,T]$; (ii) $\alpha > -e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2/\bar{\sigma}^2_{0:T}$; and (iii) $\rho > 0$.*

Proof. First, we show that

$$\alpha > -\frac{e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2}{\bar{\sigma}^2_{0:T}} \qquad \rightleftharpoons \qquad \rho > 0 \tag{38}$$

Recall that $\mathbf{x}^* = \mathbf{x}^\star + \rho(\mathbf{x}_0 - \mathbf{x}^\star)$. Thus, by Eqs. (12) and (6), we have

$$\begin{aligned}\hat{\mathbf{x}} &= \frac{\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2}\mathbf{x}^* - \frac{\sigma^2}{\bar{\sigma}^2_{0:T} - \sigma^2}\bar{\boldsymbol{\mu}}_{0:T} = -\frac{\sigma^2 e^{-\bar{\theta}_{0:T}}}{\bar{\sigma}^2_{0:T} - \sigma^2}\mathbf{x}_0 + \frac{\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2}\mathbf{x}^* - \frac{\sigma^2\left(1 - e^{-\bar{\theta}_{0:T}}\right)}{\bar{\sigma}^2_{0:T} - \sigma^2}\boldsymbol{\mu} \\ &= \left(\frac{\rho\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2} - \frac{\sigma^2 e^{-\bar{\theta}_{0:T}}}{\bar{\sigma}^2_{0:T} - \sigma^2}\right)\mathbf{x}_0 + \frac{(1-\rho)\,\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2}\mathbf{x}^\star - \frac{\sigma^2\left(1 - e^{-\bar{\theta}_{0:T}}\right)}{\bar{\sigma}^2_{0:T} - \sigma^2}\boldsymbol{\mu}.\end{aligned}$$

Combined with $\hat{\mathbf{x}} = \alpha\mathbf{x}_0 + \beta\mathbf{x}^\star + \gamma\boldsymbol{\mu}$, we have

$$\alpha = \frac{\rho\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2} - \frac{\sigma^2 e^{-\bar{\theta}_{0:T}}}{\bar{\sigma}^2_{0:T} - \sigma^2} = \rho\frac{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T}} - \frac{e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2}{\bar{\sigma}^2_{0:T}}$$

Since $0 < \sigma^2 < \bar{\sigma}^2_{0:T}$, Eq. (38) is immediate.

Next, we show that

$$\alpha > -\frac{e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2}{\bar{\sigma}^2_{0:T}} \qquad \rightleftharpoons \qquad \phi_t \neq 0 \quad \forall\, t \in [0,T]. \tag{39}$$

Then the lemma is immediate. Recall that

$$\phi_t = e^{-\bar{\theta}_{0:t}}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right) + \alpha e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right).$$

Define

$$\alpha_t = -\frac{e^{-\bar{\theta}_{0:t}}}{e^{\bar{\theta}_{t:T}}}\frac{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}}{\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}} = -e^{-\bar{\theta}_{0:T}}\frac{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}}{\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}}. \tag{40}$$

First, we show that $\phi_t \neq 0$ for all $t \in [0,T]$ implies

$$\alpha > \frac{-e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2}{\bar{\sigma}^2_{0:T}}. \tag{41}$$

Assume for contradiction that $\phi_t \neq 0$ for all $t \in [0,T]$ and

$$\alpha \leq \frac{-e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2}{\bar{\sigma}^2_{0:T}}.$$

Since $e^{\bar{\theta}_{0:T}} \geq 0$ and $\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{T:T} > 0$, we have $\phi_T \leq 0$. Note that

$$\phi_0 = e^{-\bar{\theta}_{0:0}}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}\right) = \hat{\sigma}^2 + \bar{\sigma}^2_{0:T} > 0, \;\; \phi_T = e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2 + \alpha\bar{\sigma}^2_{0:T} \leq 0.$$

Moreover, by Eq. (6), $\bar{\theta}_{0:t}$, $\bar{\theta}_{t:T}$, and $\bar{\sigma}^2_{t:T}$ are continuous functions of t. Therefore, ϕ_t is also continuous in t. Since $\phi_0 > 0$ and $\phi_T \leq 0$, the Intermediate Value Theorem implies that there exists some $t^* \in (0,T]$ such that $\phi_{t^*} = 0$, which leads to a contradiction.

Next, we show that Eq. (41) implies that $\phi_t \neq 0$ for all $t \in [0, T]$. Assume that Eq. (41) holds. Recall the definition of α_t in Eq. (40). For each fixed $t \in (0, T]$, we have $\phi_t > 0$ if and only if $\alpha > \alpha_t$. Furthermore, since $\bar{\sigma}^2_{t:T}$ is monotonically decreasing in t, and since

$$e^{-\bar{\theta}_{0:T}} > 0, \qquad \hat{\sigma}^2 > 0, \qquad \bar{\sigma}^2_{t:T} \geq 0,$$

it follows that the threshold α_t is monotonically increasing as a function of t. Hence, by Eq. (41) we obtain

$$\alpha_t \leq \alpha_T = \frac{-e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2}{\bar{\sigma}^2_{0:T}} < \alpha \qquad \forall t \in (0, T]. \tag{42}$$

Thus, $\phi_t > 0$ for all $t \in (0, T]$. Recall that $\phi_0 > 0$. We obtain that $\phi_t > 0$ for all $t \in [0, T]$. In summary, Eq. (39) is proved. The lemma is established. □

The above lemma characterizes the admissible range of α under which the reformulated dynamics remain well-defined. In particular, under the relaxed parametrization $\mathbf{x}^* = \mathbf{x}^\star + \rho(\mathbf{x}_0 - \mathbf{x}^\star)$, the condition $\phi_t \neq 0$ on $[0, T]$ is equivalent to $\rho > 0$. Hence, for the $\mathbf{x}_0$-free reformulated dynamics, the auxiliary terminal center $\mathbf{x}^*$ cannot be chosen to coincide exactly with the desired target $\mathbf{x}^\star$. Instead, a small but strictly positive relaxation is required, so that $\mathbf{x}^*$ remains close to $\mathbf{x}^\star$ while the coefficient ϕ_t stays nonzero throughout the time interval. The ρ-relaxed form is used here only to illustrate that a nonsingular $\mathbf{x}_0$-free reformulation requires a nonzero coefficient of $\mathbf{x}_0$ in the terminal mean. SDDBMs allow more general terminal centers subject to the admissible parameter conditions.

E.3. Variance Matching and Diffusion Coefficient Calibration

It remains to specify the diffusion coefficient η_t in Eq. (37). This coefficient cannot be chosen independently of the drift. Indeed, for a linear SDE of the form $\mathrm{d}\mathbf{x}_t = (f_t\mathbf{x}_t + \mathbf{b}_t)\,\mathrm{d}t + \eta_t \mathrm{d}\mathbf{w}_t$, the marginal covariance evolves according to $\frac{\mathrm{d}}{\mathrm{d}t}\mathrm{Cov}(\mathbf{x}_t) = 2f_t\,\mathrm{Cov}(\mathbf{x}_t) + \eta_t^2\mathbf{I}$. Therefore, once the drift coefficient is changed from that in Eq. (24) to the coefficient f_t in Eq. (20), the stochastic perturbation must be recalibrated so that the reformulated process preserves the target marginal distribution in Proposition 3.4. This leads to the following result.

Proposition 3. *Assume that the drift in Eq. (37) has been chosen so that the mean trajectory matches that of the h-transformed marginal distribution in Proposition 3.4. Then the time marginals are consistent if the diffusion coefficient η_t is chosen so that η_t^2 is given by Eq. (51).*

Proof. Consider the process governed by Eq. (37). Since the drift is affine in $\mathbf{x}_t$, the process is Gaussian at each time t when initialized from a deterministic $\mathbf{x}_0$. Therefore, its time marginal is fully determined by its mean and covariance. By construction, the drift in Eq. (37) matches the mean trajectory of the h-transformed marginal distribution. It remains to choose η_t so that the covariance also matches the target covariance.

Eq. (37) implies

$$\mathrm{d}\mathbf{x}_t = [f_t\mathbf{x}_t + \mathbf{b}_t]\,\mathrm{d}t + \eta_t \mathrm{d}\mathbf{w}_t, \tag{43}$$

where $\mathbf{b}_t$ is deterministic and does not affect the covariance. Thus, if $\mathrm{Cov}(\mathbf{x}_t) = V_t\mathbf{I}$, then V_t satisfies

$$\frac{\mathrm{d}V_t}{\mathrm{d}t} = 2f_tV_t + \eta_t^2. \tag{44}$$

Recall from Eqs. (12) and (19) that

$$\varphi_t = \sigma^2\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right) + \bar{\sigma}^2_{0:T}\bar{\sigma}^2_{t:T}, \qquad \hat{\sigma}^2 = \frac{\sigma^2\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2}. \tag{45}$$

Hence, the marginal covariance of the h-transformed marginal distribution in Proposition 3.4 is

$$\frac{\bar{\sigma}^2_{0:t}(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}\mathbf{I} = \frac{\bar{\sigma}^2_{0:t}\varphi_t}{\bar{\sigma}^4_{0:T}}\mathbf{I}. \tag{46}$$

Define

$$V_t^* := \frac{\bar{\sigma}^2_{0:t}\varphi_t}{\bar{\sigma}^4_{0:T}}, \qquad \eta_t^2 := \frac{\mathrm{d}V_t^*}{\mathrm{d}t} - 2f_tV_t^*. \tag{47}$$

We have $V_0^* = 0$. Moreover, we also have $V_0 = 0$ by $\mathrm{Cov}(\mathbf{x}_0) = \mathbf{0} = V_0\mathbf{I}$. Hence, Eqs. (44) and (47) imply $V_t = V_t^*$. In addition, for the GOU variance, we have

$$\frac{\mathrm{d}}{\mathrm{d}t}\bar{\sigma}^2_{t:T} = -g_t^2 e^{-2\bar{\theta}_{t:T}}, \qquad \frac{\mathrm{d}}{\mathrm{d}t}\bar{\sigma}^2_{0:t} = g_t^2 - 2\theta_t\bar{\sigma}^2_{0:t}. \tag{48}$$

Thus,

$$\frac{\mathrm{d}\varphi_t}{\mathrm{d}t} = g_t^2 e^{-2\bar{\theta}_{t:T}}\left(\sigma^2 - \bar{\sigma}^2_{0:T}\right). \tag{49}$$

Hence

$$\frac{\mathrm{d}V_t^*}{\mathrm{d}t} = \frac{1}{\bar{\sigma}^4_{0:T}}\left[\left(g_t^2 - 2\theta_t\bar{\sigma}^2_{0:t}\right)\varphi_t + \bar{\sigma}^2_{0:t}g_t^2 e^{-2\bar{\theta}_{t:T}}\left(\sigma^2 - \bar{\sigma}^2_{0:T}\right)\right]. \tag{50}$$

Substituting Eq. (50) and Eq. (20) into Eq. (47), we obtain

$$\eta_t^2 = \frac{g_t^2}{\bar{\sigma}^4_{0:T}}\left\{\varphi_t\left[1 - \frac{2e^{-\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t}\left(\alpha - e^{-\bar{\theta}_{0:T}}\right)}{\phi_t}\right] + \bar{\sigma}^2_{0:t}e^{-2\bar{\theta}_{t:T}}\left(\sigma^2 - \bar{\sigma}^2_{0:T}\right)\right\}. \tag{51}$$

Consequently, the reformulated process has the same covariance as the target h-transformed marginal distribution. Since the drift has already been chosen to match the target mean and both marginals are Gaussian, the time marginals are consistent. □

This choice of η_t ensures that the stochastic perturbation in Eq. (37) compensates for the modified linear drift and preserves the marginal covariance specified by Proposition 3.4. In this sense, Eq. (37) should be interpreted as a marginally consistent reformulation of Eq. (24), rather than as a pathwise identity obtained by a direct substitution inside the stochastic term.

Eq. (51) can also be equivalently written as

$$\eta_t^2 = g_t^2\frac{e^{-\bar{\theta}_{0:T}}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right) - \alpha\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right)}{e^{-\bar{\theta}_{0:T}}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right) + \alpha\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right)}. \tag{52}$$

The condition

$$-\frac{e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2}{\bar{\sigma}^2_{0:T}} < \alpha \le \frac{e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2}{\bar{\sigma}^2_{0:T}} \tag{53}$$

ensures nonsingular drift coefficients (by Lemma 2) and $\eta_t^2 \ge 0$ for all $t \in [0, T]$.

By Propositions 1 and 3, the reverse SDE in Eq. (18) follows immediately from the frameworks established in (Song et al., 2021b; Yue et al., 2024).

F. Missing proofs and derivations

F.1. Proof of Proposition 3.1

Proposition 3.1. *Consider the affine diffusion dynamics governed by the SDE* $\mathrm{d}\mathbf{x}_t = (A_t\mathbf{x}_t + \mathbf{b}_t)\mathrm{d}t + g_t\mathrm{d}\mathbf{W}_t$ *on a finite time interval* $[0, T]$*, where* $A : [0, T] \to \mathbb{R}^{d\times d}$*,* $\mathbf{b} : [0, T] \to \mathbb{R}^d$*, and* $g : [0, T] \to \mathbb{R}$ *are deterministic, time-dependent coefficients. Notably, for any time* t*, these parameters are strictly independent of the initial state* $\mathbf{x}_0$*, the current state* $\mathbf{x}_t$*, and the driving Brownian motion. Assume that*

$$\sup_{0\le t\le T}\|A_t\| < \infty, \qquad \sup_{0\le t\le T}\|\mathbf{b}_t\| < \infty, \qquad 0 < \int_0^T g_t^2\mathrm{d}t < \infty. \tag{9}$$

Let $\Phi(t, s)$ *be the fundamental matrix associated with* A_t*, defined by* $\Phi(s, s) = \mathbf{I}$*,* $\frac{\partial}{\partial t}\Phi(t, s) = A_t\Phi(t, s)$*. Then the terminal law of* $\mathbf{x}_T$ *is a non-degenerate Gaussian:* $\mathbf{x}_T \mid \mathbf{x}_0 \sim \mathcal{N}\left(\mathbf{m}_T(\mathbf{x}_0), \Sigma_T\right)$ *where* $\mathbf{m}_T(\mathbf{x}_0) = \Phi(T, 0)\mathbf{x}_0 + \int_0^T \Phi(T, s)\mathbf{b}_s\,\mathrm{d}s, \Sigma_T = \int_0^T g_s^2\Phi(T, s)\Phi(T, s)^\top\,\mathrm{d}s$*. Here,* $\Phi(T, 0)$ *is invertible and* Σ_T *is positive definite. Consequently,* $p(\mathbf{x}, T \mid \mathbf{x}_0, 0)$ *is a smooth, continuous, and bounded density. Moreover,* $\mathbf{m}_T(\mathbf{x}_0) \ne \mathbf{m}_T(\mathbf{y}_0)$ *for any* $\mathbf{x}_0 \ne \mathbf{y}_0$*.*

Proof. Let $R_t := \Phi(t, 0)$. Then R_t satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t} R_t = A_t R_t, \qquad R_0 = \mathbf{I}.$$

Since A_t is uniformly bounded on the finite interval $[0, T]$, the matrix R_t is invertible for every $t \in [0, T]$. Moreover,

$$\frac{\mathrm{d}}{\mathrm{d}t} R_t^{-1} = -R_t^{-1} A_t.$$

Define $\mathbf{y}_t := R_t^{-1} \mathbf{x}_t$. Using Itô's product rule and the SDE $\mathrm{d}\mathbf{x}_t = (A_t \mathbf{x}_t + \mathbf{b}_t)\mathrm{d}t + g_t \mathrm{d}\mathbf{W}_t$,

$$\begin{aligned}
\mathrm{d}\mathbf{y}_t &= \mathrm{d}(R_t^{-1} \mathbf{x}_t) \\
&= \left(\frac{\mathrm{d}}{\mathrm{d}t} R_t^{-1} \right) \mathbf{x}_t \mathrm{d}t + R_t^{-1} \mathrm{d}\mathbf{x}_t \\
&= -R_t^{-1} A_t \mathbf{x}_t \mathrm{d}t + R_t^{-1} \left[(A_t \mathbf{x}_t + \mathbf{b}_t)\mathrm{d}t + g_t \mathrm{d}\mathbf{W}_t \right] \\
&= R_t^{-1} \mathbf{b}_t \mathrm{d}t + R_t^{-1} g_t \mathrm{d}\mathbf{W}_t.
\end{aligned}$$

Integrating from 0 to T gives

$$\mathbf{y}_T = \mathbf{x}_0 + \int_0^T R_s^{-1} \mathbf{b}_s \, \mathrm{d}s + \int_0^T R_s^{-1} g_s \, \mathrm{d}\mathbf{W}_s.$$

Multiplying both sides by R_T, and using $R_T R_s^{-1} = \Phi(T, s)$, we obtain the variation-of-constants representation

$$\mathbf{x}_T = \Phi(T, 0)\mathbf{x}_0 + \int_0^T \Phi(T, s) \mathbf{b}_s \, \mathrm{d}s + \int_0^T \Phi(T, s) g_s \, \mathrm{d}\mathbf{W}_s. \tag{54}$$

The first two terms in Eq. (54) are deterministic given a fixed $\mathbf{x}_0$. The last term is a Gaussian random vector, since it is an Itô integral with deterministic integrand. Therefore, $\mathbf{x}_T \mid \mathbf{x}_0$ is Gaussian with mean

$$\mathbf{m}_T(\mathbf{x}_0) = \Phi(T, 0)\mathbf{x}_0 + \int_0^T \Phi(T, s) \mathbf{b}_s \, \mathrm{d}s$$

and covariance

$$\Sigma_T = \int_0^T g_s^2 \Phi(T, s) \Phi(T, s)^\top \, \mathrm{d}s.$$

It remains to show that Σ_T is positive definite. Let

$$M_A := \sup_{0 \le t \le T} \|A_t\| < \infty.$$

By Gronwall's inequality, for $0 \le s \le T$,

$$\|\Phi(T, s)\| \le e^{M_A (T - s)} \le e^{M_A T}.$$

Similarly, since $\Phi(T, s)$ is invertible and $\Phi(T, s)^{-1} = \Phi(s, T)$, we also have

$$\|\Phi(T, s)^{-1}\| = \|\Phi(s, T)\| \le e^{M_A (T - s)} \le e^{M_A T}.$$

Hence, for any $\mathbf{v} \in \mathbb{R}^d$,

$$\|\Phi(T, s)^\top \mathbf{v}\| \ge e^{-M_A T} \|\mathbf{v}\|.$$

Therefore,

$$\begin{aligned}
\mathbf{v}^\top \Sigma_T \mathbf{v} &= \int_0^T g_s^2 \mathbf{v}^\top \Phi(T, s) \Phi(T, s)^\top \mathbf{v} \, \mathrm{d}s \\
&= \int_0^T g_s^2 \|\Phi(T, s)^\top \mathbf{v}\|^2 \, \mathrm{d}s \\
&\ge e^{-2 M_A T} \left(\int_0^T g_s^2 \, \mathrm{d}s \right) \|\mathbf{v}\|^2.
\end{aligned}$$

By Eq. (9), we obtain

$$\mathbf{v}^\top \Sigma_T \mathbf{v} > 0, \qquad \forall\, \mathbf{v} \neq 0.$$

Thus, $\Sigma_T \succ 0$. Since Σ_T is positive definite, the terminal Gaussian is non-degenerate and has density

$$p(\mathbf{x}, T \mid \mathbf{x}_0, 0) = \frac{1}{(2\pi)^{d/2}|\Sigma_T|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{m}_T(\mathbf{x}_0))^\top \Sigma_T^{-1}(\mathbf{x} - \mathbf{m}_T(\mathbf{x}_0))\right).$$

Thus, as a function of $\mathbf{x}$, this density is smooth, i.e., it has continuous derivatives of all orders. In particular, it is continuous. Moreover, since

$$p(\mathbf{x}, T \mid \mathbf{x}_0, 0) \leq \frac{1}{(2\pi)^{d/2}|\Sigma_T|^{1/2}} < \infty,$$

the density is bounded.

Finally, we prove that the terminal density depends on the source state. For two initial states $\mathbf{x}_0, \mathbf{y}_0$, the difference between the terminal means is

$$\mathbf{m}_T(\mathbf{x}_0) - \mathbf{m}_T(\mathbf{y}_0) = \Phi(T, 0)(\mathbf{x}_0 - \mathbf{y}_0).$$

Since $\Phi(T, 0)$ is invertible, we have $\det \Phi(T, 0) \neq 0$. Therefore, if $\mathbf{x}_0 \neq \mathbf{y}_0$, then

$$\Phi(T, 0)(\mathbf{x}_0 - \mathbf{y}_0) \neq 0,$$

and hence $\mathbf{m}_T(\mathbf{x}_0) \neq \mathbf{m}_T(\mathbf{y}_0)$. This completes the proof. □

F.2. Proof of Theorem 3.2

Theorem 3.2. *Consider the GOU process initialized at $\mathbf{x}_0$, with terminal transition density $p(\mathbf{x}, T \mid \mathbf{x}_0, 0) = \mathcal{N}\left(\mathbf{x}; \bar{\boldsymbol{\mu}}_{0:T}, \bar{\sigma}^2_{0:T}\mathbf{I}\right)$, where $\bar{\boldsymbol{\mu}}_{0:T}$ and $\bar{\sigma}^2_{0:T}$ denote the conditional mean and variance of the state at time T given $\mathbf{x}_0$, respectively. Let q be a nonnegative probability density and define*

$$h(\mathbf{x}, t) = \int p(\mathbf{z}, T \mid \mathbf{x}, t) q(\mathbf{z})\, \mathrm{d}\mathbf{z}. \tag{11}$$

Assume that the terminal marginal under this transformed measure is $p_h(\mathbf{x}, T \mid \mathbf{x}_0, 0) = \mathcal{N}(\mathbf{x}; \mathbf{x}^*, \sigma^2\mathbf{I})$ for some $0 < \sigma^2 < \bar{\sigma}^2_{0:T}$. Then q is uniquely determined up to almost-everywhere equality and is given by $q(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \hat{\mathbf{x}}, \hat{\sigma}^2\mathbf{I})$, where

$$\hat{\mathbf{x}} = \frac{\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2}\mathbf{x}^* - \frac{\sigma^2}{\bar{\sigma}^2_{0:T} - \sigma^2}\bar{\boldsymbol{\mu}}_{0:T}, \qquad \hat{\sigma}^2 = \frac{\sigma^2\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2}. \tag{12}$$

Proof. By assumption,

$$p_h(\mathbf{x}, T \mid \mathbf{x}_0, 0) = \mathcal{N}(\mathbf{x}; \mathbf{x}^*, \sigma^2\mathbf{I}), \qquad p(\mathbf{x}, T \mid \mathbf{x}_0, 0) = \mathcal{N}(\mathbf{x}; \bar{\boldsymbol{\mu}}_{0:T}, \bar{\sigma}^2_{0:T}\mathbf{I}).$$

Since $p(\mathbf{x}, T \mid \mathbf{x}_0, 0) > 0$ for all $\mathbf{x}$, Eq. (4) implies

$$\mathcal{N}(\mathbf{x}; \mathbf{x}^*, \sigma^2\mathbf{I}) = \frac{\mathcal{N}(\mathbf{x}; \bar{\boldsymbol{\mu}}_{0:T}, \bar{\sigma}^2_{0:T}\mathbf{I}) q(\mathbf{x})}{h(\mathbf{x}_0, 0)}.$$

Equivalently,

$$q(\mathbf{x}) = h(\mathbf{x}_0, 0)\frac{\mathcal{N}(\mathbf{x}; \mathbf{x}^*, \sigma^2\mathbf{I})}{\mathcal{N}(\mathbf{x}; \bar{\boldsymbol{\mu}}_{0:T}, \bar{\sigma}^2_{0:T}\mathbf{I})}.$$

Hence,

$$q(\mathbf{x}) \propto \exp\left\{-\frac{1}{2}\left[\frac{\|\mathbf{x} - \mathbf{x}^*\|^2}{\sigma^2} - \frac{\|\mathbf{x} - \bar{\boldsymbol{\mu}}_{0:T}\|^2}{\bar{\sigma}^2_{0:T}}\right]\right\}.$$

By

$$0 < \sigma^2 < \bar{\sigma}^2_{0:T},$$

we obtain

$$\frac{1}{\sigma^2} - \frac{1}{\bar{\sigma}^2_{0:T}} > 0.$$

Hence, since q is a nonnegative probability density, we obtain

$$q(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \hat{\mathbf{x}}, \hat{\sigma}^2 \mathbf{I}),$$

where

$$\hat{\mathbf{x}} = \hat{\sigma}^2 \left(\frac{\mathbf{x}^*}{\sigma^2} - \frac{\bar{\boldsymbol{\mu}}_{0:T}}{\bar{\sigma}^2_{0:T}} \right) = \frac{\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2} \mathbf{x}^* - \frac{\sigma^2}{\bar{\sigma}^2_{0:T} - \sigma^2} \bar{\boldsymbol{\mu}}_{0:T},$$
$$\hat{\sigma}^2 = \left(\frac{1}{\sigma^2} - \frac{1}{\bar{\sigma}^2_{0:T}} \right)^{-1} = \frac{\sigma^2 \bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2}.$$

The theorem is proved. □

F.3. Proof of Theorem 3.3

Theorem 3.3. *Consider the transition kernel of the GOU process, given by*

$$p(\mathbf{z}, T | \mathbf{x}_t, t) = \mathcal{N}(\mathbf{z}; \bar{\boldsymbol{\mu}}_{t:T}, \bar{\sigma}^2_{t:T} \mathbf{I}), \tag{13}$$

where $\bar{\boldsymbol{\mu}}_{t:T}$ and $\bar{\sigma}^2_{t:T}$ denote the conditional mean and variance of the state at time T given $\mathbf{x}_t$, respectively. Let $q(\mathbf{z}) := \mathcal{N}(\mathbf{z}; \hat{\mathbf{x}}, \hat{\sigma}^2 \mathbf{I})$ be a Gaussian density with parameters $\hat{\mathbf{x}} \in \mathbb{R}^d$ and $\hat{\sigma} > 0$. We define the function $h(\mathbf{x}_t, t)$ via the integral $h(\mathbf{x}_t, t) := \int p(\mathbf{z}, T | \mathbf{x}_t, t) q(\mathbf{z}) \, \mathrm{d}\mathbf{z}$. For any $t \in [0, T]$, $h(\mathbf{x}_t, t)$ admits the closed-form expression:

$$h(\mathbf{x}_t, t) = \frac{1}{[2\pi(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})]^{d/2}} \exp\left[-\frac{\|\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}\|^2}{2(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})} \right]. \tag{14}$$

Furthermore, for any $t \in [0, T)$, $h(\mathbf{x}_t, t)$ satisfies the Kolmogorov Backward Equation.

Proof. Using the standard product identity for Gaussian densities (Petersen & Pedersen, 2012),

$$\int_{\mathbb{R}^d} \mathcal{N}(\mathbf{z}; \mathbf{m}_1, \Sigma_1) \mathcal{N}(\mathbf{z}; \mathbf{m}_2, \Sigma_2) \, \mathrm{d}\mathbf{z} = \mathcal{N}(\mathbf{m}_1; \mathbf{m}_2, \Sigma_1 + \Sigma_2), \tag{55}$$

we obtain

$$\begin{aligned} h(\mathbf{x}_t, t) &= \int \mathcal{N}(\mathbf{z}; \bar{\boldsymbol{\mu}}_{t:T}, \bar{\sigma}^2_{t:T} \mathbf{I}) \mathcal{N}(\mathbf{z}; \hat{\mathbf{x}}, \hat{\sigma}^2 \mathbf{I}) \, \mathrm{d}\mathbf{z} \\ &= \mathcal{N}(\bar{\boldsymbol{\mu}}_{t:T}; \hat{\mathbf{x}}, (\bar{\sigma}^2_{t:T} + \hat{\sigma}^2) \mathbf{I}) \\ &= \frac{1}{[2\pi(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})]^{d/2}} \exp\left[-\frac{\|\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}\|^2}{2(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})} \right]. \end{aligned} \tag{56}$$

Next, we verify that $h(\mathbf{x}_t, t)$ satisfies the KBE. We first compute $\nabla_{\mathbf{x}_t} h(\mathbf{x}_t, t)$ and $\nabla^2_{\mathbf{x}_t} h(\mathbf{x}_t, t)$:

$$\begin{aligned} \nabla_{\mathbf{x}_t} h(\mathbf{x}_t, t) &= -h(\mathbf{x}_t, t) \frac{1}{2(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})} \frac{\partial \|\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}\|^2}{\partial \mathbf{x}_t} \\ &= -h(\mathbf{x}_t, t) \left(\frac{\partial \bar{\boldsymbol{\mu}}_{t:T}}{\partial \mathbf{x}_t} \right)^\top \frac{\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}} \\ &= -h(\mathbf{x}_t, t) e^{-\bar{\theta}_{t:T}} \frac{\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}}, \\ \nabla^2_{\mathbf{x}_t} h(\mathbf{x}_t, t) &= \nabla_{\mathbf{x}_t} \left(\nabla_{\mathbf{x}_t} h(\mathbf{x}_t, t) \right) \\ &= -\nabla_{\mathbf{x}_t} \left[h(\mathbf{x}_t, t) e^{-\bar{\theta}_{t:T}} \frac{\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}} \right] \\ &= -h(\mathbf{x}_t, t) \frac{e^{-2\bar{\theta}_{t:T}} \mathbf{I}}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}} + h(\mathbf{x}_t, t) e^{-2\bar{\theta}_{t:T}} \frac{(\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}})(\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}})^\top}{(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})^2} \\ &= h(\mathbf{x}_t, t) e^{-2\bar{\theta}_{t:T}} \left[\frac{(\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}})(\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}})^\top}{(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})^2} - \frac{\mathbf{I}}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}} \right]. \end{aligned}$$

Hence, we obtain

$$\Delta_{\mathbf{x}_t} h(\mathbf{x}_t, t) = h(\mathbf{x}_t, t) e^{-2\bar{\theta}_{t:T}} \left[\frac{\|\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}\|^2}{(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})^2} - \frac{d}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}} \right].$$

Moreover,

$$\begin{aligned}
\frac{\partial h(\mathbf{x}_t, t)}{\partial t} =& -\frac{d}{2}(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})^{-1} \frac{\partial \bar{\sigma}^2_{t:T}}{\partial t} h(\mathbf{x}_t, t) - \frac{1}{2} h(\mathbf{x}_t, t) \frac{\partial}{\partial t} \left[\frac{\|\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}\|^2}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}} \right] \\
=& h(\mathbf{x}_t, t) \frac{g_t^2 e^{-2\bar{\theta}_{t:T}} d}{2(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})} - \frac{1}{2} h(\mathbf{x}_t, t) \left[\frac{2(\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}})^\top}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}} \frac{\partial \bar{\boldsymbol{\mu}}_{t:T}}{\partial t} - \frac{\|\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}\|^2}{(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})^2} \frac{\partial \bar{\sigma}^2_{t:T}}{\partial t} \right] \\
=& h(\mathbf{x}_t, t) \left[\frac{g_t^2 e^{-2\bar{\theta}_{t:T}} d}{2(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})} - \frac{(\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}})^\top}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}} \frac{\partial \bar{\boldsymbol{\mu}}_{t:T}}{\partial t} + \frac{\|\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}\|^2}{2(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})^2} \frac{\partial \bar{\sigma}^2_{t:T}}{\partial t} \right] \\
=& h(\mathbf{x}_t, t) \left[\frac{g_t^2 e^{-2\bar{\theta}_{t:T}} d}{2(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})} - \frac{(\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}})^\top (\mathbf{x}_t - \boldsymbol{\mu})}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}} e^{-\bar{\theta}_{t:T}} \theta_t - \frac{\|\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}\|^2}{2(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})^2} g_t^2 e^{-2\bar{\theta}_{t:T}} \right].
\end{aligned}$$

Therefore,

$$\begin{aligned}
&\theta_t(\boldsymbol{\mu} - \mathbf{x}_t) \cdot \nabla_{\mathbf{x}_t} h(\mathbf{x}_t, t) + \frac{1}{2} g_t^2 \Delta_{\mathbf{x}_t} h(\mathbf{x}_t, t) \\
=& -\theta_t h(\mathbf{x}_t, t) \frac{(\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}})^\top (\boldsymbol{\mu} - \mathbf{x}_t)}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}} e^{-\bar{\theta}_{t:T}} + \frac{1}{2} g_t^2 h(\mathbf{x}_t, t) e^{-2\bar{\theta}_{t:T}} \left[\frac{\|\bar{\boldsymbol{\mu}}_{t:T} - \hat{\mathbf{x}}\|^2}{(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})^2} - \frac{d}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}} \right] \\
=& -\frac{\partial h(\mathbf{x}_t, t)}{\partial t}.
\end{aligned}$$

Thus, the h-function satisfies the KBE. □

F.4. Proof of Proposition 3.4

Proposition 3.4. *Consider the stochastic process governed by Eq. (15). For $0 < t \leq T$, its time-t marginal density under the h-transformed path measure is*

$$p_h(\mathbf{x}_t, t | \mathbf{x}_0, 0) = \mathcal{N}\left(\mathbf{x}_t; \bar{\boldsymbol{\mu}}_{0:t} + \frac{e^{-\bar{\theta}_{t:T}} \bar{\sigma}^2_{0:t}}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}} (\hat{\mathbf{x}} - \bar{\boldsymbol{\mu}}_{0:T}), \frac{\bar{\sigma}^2_{0:t} (\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}} \mathbf{I} \right). \tag{17}$$

Proof. By applying the standard marginalization and conditioning properties of Gaussian distributions (Petersen & Pedersen, 2012), we can rewrite the product of the marginal and conditional densities. Specifically, assuming A, S, and R are positive scalars, for isotropic Gaussian variables

$$\mathbf{x} \sim \mathcal{N}(\mathbf{m}, S\mathbf{I}), \qquad \mathbf{y} \mid \mathbf{x} \sim \mathcal{N}(A\mathbf{x} + \mathbf{c}, R\mathbf{I}),$$

the joint density $p(\mathbf{x}, \mathbf{y})$ can be factored as $p(\mathbf{y})p(\mathbf{x} \mid \mathbf{y})$, which yields the following identity:

$$\begin{aligned}
&\mathcal{N}(\mathbf{x}; \mathbf{m}, S\mathbf{I}) \mathcal{N}(\mathbf{y}; A\mathbf{x} + \mathbf{c}, R\mathbf{I}) \\
=& \mathcal{N}(\mathbf{y}; A\mathbf{m} + \mathbf{c}, (A^2 S + R)\mathbf{I}) \mathcal{N}\left(\mathbf{x}; \mathbf{m} + \frac{SA}{A^2 S + R} (\mathbf{y} - A\mathbf{m} - \mathbf{c}), \frac{SR}{A^2 S + R} \mathbf{I} \right).
\end{aligned}$$

Within our framework, we instantiate the parameters as follows:

$$\mathbf{x} = \mathbf{x}_t, \quad \mathbf{y} = \hat{\mathbf{x}}, \quad \mathbf{m} = \bar{\boldsymbol{\mu}}_{0:t}, \quad S = \bar{\sigma}^2_{0:t}, \quad A = e^{-\bar{\theta}_{t:T}}, \quad \mathbf{c} = \boldsymbol{\mu}(1 - e^{-\bar{\theta}_{t:T}}), \quad R = \hat{\sigma}^2 + \bar{\sigma}^2_{t:T}.$$

Substituting these into the aforementioned identities yields the desired marginal parameters:

$$A\mathbf{m} + \mathbf{c} = \bar{\boldsymbol{\mu}}_{0:T},$$

and

$$A^2 S + R = \hat{\sigma}^2 + \bar{\sigma}^2_{0:T}.$$

Hence, by Eqs. (13) and (14),

$$p(\mathbf{x}_t, t \mid \mathbf{x}_0, 0)h(\mathbf{x}_t, t) = \mathcal{N}(\hat{\mathbf{x}}; \bar{\boldsymbol{\mu}}_{0:T}, \left(\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}\right)\mathbf{I})\mathcal{N}\left(\mathbf{x}_t; \bar{\boldsymbol{\mu}}_{0:t} + \frac{e^{-\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t}}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}(\hat{\mathbf{x}} - \bar{\boldsymbol{\mu}}_{0:T}), \frac{\bar{\sigma}^2_{0:t}(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}\mathbf{I}\right)$$
$$= h(\mathbf{x}_0, 0)\mathcal{N}\left(\mathbf{x}_t; \bar{\boldsymbol{\mu}}_{0:t} + \frac{e^{-\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t}}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}(\hat{\mathbf{x}} - \bar{\boldsymbol{\mu}}_{0:T}), \frac{\bar{\sigma}^2_{0:t}(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}\mathbf{I}\right).$$

By Eq. (2), we obtain

$$p_h(\mathbf{x}_t, t \mid \mathbf{x}_0, 0) = \frac{p(\mathbf{x}_t, t|\mathbf{x}_0, 0)h(\mathbf{x}_t, t)}{h(\mathbf{x}_0, 0)} = \mathcal{N}\left(\mathbf{x}_t; \bar{\boldsymbol{\mu}}_{0:t} + \frac{e^{-\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t}}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}(\hat{\mathbf{x}} - \bar{\boldsymbol{\mu}}_{0:T}), \frac{\bar{\sigma}^2_{0:t}(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T})}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}\mathbf{I}\right).$$

□

F.5. Proof of Proposition 4.1

Proposition 4.1. *The proposed SDDBMs framework recovers DDBM-VP, DDBM-VE, GOUB, and UniDB as special cases under the hyperparameter choices listed in Table 1.*

Proof. We prove the proposition by substituting the corresponding hyper-parameter choices into the general reformulated SDDBM dynamics in Eq. (37). Let $\alpha = 0, \beta = 1, \gamma = 0$. Recall that

$$\phi_t = e^{-\bar{\theta}_{0:t}}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right) + \alpha e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right), \tag{57}$$

$$\hat{\sigma}^2 = \frac{\sigma^2\bar{\sigma}^2_{0:T}}{\bar{\sigma}^2_{0:T} - \sigma^2}. \tag{58}$$

Moreover, for the GOU process,

$$e^{-\bar{\theta}_{0:T}} = e^{-\bar{\theta}_{0:t}}e^{-\bar{\theta}_{t:T}}, \qquad \bar{\sigma}^2_{0:T} = e^{-2\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t} + \bar{\sigma}^2_{t:T}. \tag{59}$$

By $\alpha = 0$ and Eq. (57), we have

$$\phi_t = e^{-\bar{\theta}_{0:t}}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right). \tag{60}$$

Substituting $\alpha = 0$ and Eq. (60) into Eq. (51), we obtain

$$\eta_t^2 = \frac{g_t^2}{\bar{\sigma}^4_{0:T}}\left\{\varphi_t\left[1 + \frac{2e^{-2\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t}}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}}\right] + \bar{\sigma}^2_{0:t}e^{-2\bar{\theta}_{t:T}}\left(\sigma^2 - \bar{\sigma}^2_{0:T}\right)\right\}. \tag{61}$$

By Eq. (58), we obtain

$$\varphi_t = \frac{\bar{\sigma}^4_{0:T}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right)}{\bar{\sigma}^2_{0:T} + \hat{\sigma}^2}, \qquad \sigma^2 - \bar{\sigma}^2_{0:T} = -\frac{\bar{\sigma}^4_{0:T}}{\bar{\sigma}^2_{0:T} + \hat{\sigma}^2}. \tag{62}$$

Substituting Eq. (62) into Eq. (61) and using Eq. (59) gives

$$\begin{aligned}\eta_t^2 &= \frac{g_t^2}{\bar{\sigma}^4_{0:T}}\left[\frac{\bar{\sigma}^4_{0:T}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right)}{\bar{\sigma}^2_{0:T} + \hat{\sigma}^2}\left(1 + \frac{2e^{-2\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t}}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}}\right) - \frac{e^{-2\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t}\bar{\sigma}^4_{0:T}}{\bar{\sigma}^2_{0:T} + \hat{\sigma}^2}\right] \\ &= g_t^2\frac{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T} + e^{-2\bar{\theta}_{t:T}}\bar{\sigma}^2_{0:t}}{\bar{\sigma}^2_{0:T} + \hat{\sigma}^2} \\ &= g_t^2.\end{aligned} \tag{63}$$

Therefore, $\eta_t = g_t$ holds for all cases with $\alpha = 0$.

We now simplify the drift under the same hyper-parameter choice. By Eqs. (59) and (60), we have

$$\frac{e^{-\bar{\theta}_{0:T}}}{\phi_t} = \frac{e^{-\bar{\theta}_{t:T}}}{\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}}. \tag{64}$$

From Eqs. (20), (21), and (64), we obtain

$$f_t = -\theta_t - g_t^2 \frac{e^{-2\bar{\theta}_{t:T}}}{(\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2)}, \tag{65}$$

$$m_t = g_t^2 \frac{\beta e^{-\bar{\theta}_{t:T}}}{\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2}, \tag{66}$$

and

$$h_t = \theta_t + g_t^2 \frac{e^{-\bar{\theta}_{t:T}} \left[e^{-\bar{\theta}_{t:T}} + (\gamma - 1) \right]}{\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2}. \tag{67}$$

DDBM-VP. For DDBM-VP, we choose

$$\theta_t = \frac{1}{2} g_t^2, \qquad \boldsymbol{\mu} = \mathbf{0}, \qquad \alpha = 0, \quad \beta = 1, \quad \gamma = 0, \qquad \hat{\sigma} = 0.$$

Then

$$\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2 = \bar{\sigma}_{t:T}^2.$$

Since $\boldsymbol{\mu} = \mathbf{0}$, the term $h_t \boldsymbol{\mu}$ vanishes. Thus, by Eqs. (65) and (66), the drift in Eq. (37) becomes

$$f_t \mathbf{x}_t + m_t \mathbf{x}^\star = -\left(\frac{1}{2} g_t^2 + g_t^2 \frac{e^{-2\bar{\theta}_{t:T}}}{\bar{\sigma}_{t:T}^2} \right) \mathbf{x}_t + g_t^2 \frac{e^{-\bar{\theta}_{t:T}}}{\bar{\sigma}_{t:T}^2} \mathbf{x}^\star. \tag{68}$$

Together with $\eta_t = g_t$, this yields

$$\mathrm{d}\mathbf{x}_t = \left[-\left(\frac{1}{2} g_t^2 + g_t^2 \frac{e^{-2\bar{\theta}_{t:T}}}{\bar{\sigma}_{t:T}^2} \right) \mathbf{x}_t + g_t^2 \frac{e^{-\bar{\theta}_{t:T}}}{\bar{\sigma}_{t:T}^2} \mathbf{x}^\star \right] \mathrm{d}t + g_t d\mathbf{w}_t, \tag{69}$$

which recovers the DDBM bridge with a VP reference process.

DDBM-VE. For DDBM-VE, we choose

$$\theta_t = 0, \qquad \boldsymbol{\mu} = \mathbf{0}, \qquad \alpha = 0, \quad \beta = 1, \quad \gamma = 0, \qquad \hat{\sigma} = 0.$$

Then

$$\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2 = \bar{\sigma}_{t:T}^2.$$

Since $\theta_t = 0$, we have

$$\bar{\theta}_{t:T} = 0, \qquad e^{-\bar{\theta}_{t:T}} = 1.$$

Therefore, by Eqs. (65) and (66),

$$f_t = -g_t^2 \frac{1}{\bar{\sigma}_{t:T}^2}, \qquad m_t = g_t^2 \frac{1}{\bar{\sigma}_{t:T}^2}.$$

Again $h_t \boldsymbol{\mu} = 0$, and the drift in Eq. (37) becomes

$$f_t \mathbf{x}_t + m_t \mathbf{x}^\star = g_t^2 \frac{\mathbf{x}^\star - \mathbf{x}_t}{\bar{\sigma}_{t:T}^2}.$$

Together with $\eta_t = g_t$, we obtain

$$\mathrm{d}\mathbf{x}_t = g_t^2 \frac{\mathbf{x}^\star - \mathbf{x}_t}{\bar{\sigma}_{t:T}^2} \mathrm{d}t + g_t d\mathbf{w}_t, \tag{70}$$

which recovers the DDBM bridge with a VE reference process.

GOUB. For GOUB, we choose

$$\theta_t = \frac{1}{2\lambda^2} g_t^2, \qquad \boldsymbol{\mu} = \mathbf{x}^\star, \qquad \alpha = 0, \quad \beta = 1, \quad \gamma = 0, \qquad \hat{\sigma} = 0.$$

By $\boldsymbol{\mu} = \mathbf{x}^\star$, Eqs. (65), (66) and (67), the drift in Eq. (37) becomes

$$\begin{aligned} & f_t \mathbf{x}_t + m_t \mathbf{x}^\star + h_t \boldsymbol{\mu} = f_t \mathbf{x}_t + (m_t + h_t)\mathbf{x}^\star \\ &= -\left(\theta_t + g_t^2 \frac{e^{-2\bar{\theta}_{t:T}}}{\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2}\right)\mathbf{x}_t + \left(g_t^2 \frac{e^{-\bar{\theta}_{t:T}}}{\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2} + \theta_t + g_t^2 \frac{e^{-\bar{\theta}_{t:T}}\left(e^{-\bar{\theta}_{t:T}} - 1\right)}{\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2}\right)\mathbf{x}^\star \\ &= \left(\theta_t + g_t^2 \frac{e^{-2\bar{\theta}_{t:T}}}{\hat{\sigma}^2 + \bar{\sigma}_{t:T}^2}\right)(\mathbf{x}^\star - \mathbf{x}_t)\,. \end{aligned} \tag{71}$$

Together with $\hat{\sigma} = 0$ and $\eta_t = g_t$, we obtain

$$\mathrm{d}\mathbf{x}_t = \left(\theta_t + g_t^2 \frac{e^{-2\bar{\theta}_{t:T}}}{\bar{\sigma}_{t:T}^2}\right)(\mathbf{x}^\star - \mathbf{x}_t)\,\mathrm{d}t + g_t d\mathbf{w}_t, \tag{72}$$

which is the GOUB bridge dynamics.

UniDB. For UniDB, we choose

$$\theta_t = \frac{1}{2\lambda^2} g_t^2, \qquad \boldsymbol{\mu} = \mathbf{x}^\star, \qquad \alpha = 0, \quad \beta = 1, \quad \gamma = 0, \qquad \hat{\sigma}^2 = \kappa^{-1},$$

where κ denotes the UniDB terminal penalty coefficient. Since $\boldsymbol{\mu} = \mathbf{x}^\star$, we also have Eq. (71). Together with $\hat{\sigma}^2 = \kappa^{-1}$ and $\eta_t = g_t$, we obtain

$$\mathrm{d}\mathbf{x}_t = \left(\theta_t + g_t^2 \frac{e^{-2\bar{\theta}_{t:T}}}{\kappa^{-1} + \bar{\sigma}_{t:T}^2}\right)(\mathbf{x}^\star - \mathbf{x}_t)\,\mathrm{d}t + g_t d\mathbf{w}_t, \tag{73}$$

which recovers the UniDB dynamics. □

F.6. Derivation of the training objective

In this section, we derive the training objective for the proposed SDDBMs. The derivation follows the conditional score matching formulation (Song et al., 2021b; Vincent, 2011), based on the soft h-transformed forward bridge kernel $p_h(\mathbf{x}_t, t \mid \mathbf{x}_0, 0)$. We first specify the forward marginals and a marginally consistent Gaussian one-step transition, derive the corresponding posterior, and then formulate the training objective using a noise-parameterized reverse transition.

For notational simplicity, when no ambiguity arises, we omit the explicit time indices in the transition densities; for example, we write

$$p_h(\mathbf{x}_t|\mathbf{x}_0) := p_h(\mathbf{x}_t, t|\mathbf{x}_0, 0),$$

and similarly abbreviate

$$\begin{aligned} p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t) &:= p_\theta(\mathbf{x}_{t-1}, t-1 \mid \mathbf{x}_t, t), \\ p_h(\mathbf{x}_{t-1} \mid \mathbf{x}_0, \mathbf{x}_t) &:= p_h(\mathbf{x}_{t-1}, t-1 \mid \mathbf{x}_0, 0, \mathbf{x}_t, t). \end{aligned}$$

From Proposition 3.4, the forward marginal of the proposed bridge is

$$p_h(\mathbf{x}_t \mid \mathbf{x}_0) = \mathcal{N}\left(\mathbf{x}_t; \bar{\boldsymbol{\mu}}'_{0:t}, \bar{\sigma}'^2_{0:t}\mathbf{I}\right), \tag{74}$$

where

$$\bar{\boldsymbol{\mu}}'_{0:t} = a_t \mathbf{x}_0 + b_t \mathbf{x}^\star + c_t \boldsymbol{\mu}, \qquad \bar{\sigma}'^2_{0:t} = \bar{\sigma}^2_{0:t}(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}) / \left(\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}\right). \tag{75}$$

The coefficients a_t, b_t, and c_t correspond to $\mathbf{x}_0$, $\mathbf{x}^\star$, and $\boldsymbol{\mu}$, respectively, in the mean trajectory derived in Appendix E.1. Their closed-form expressions are

$$a_t = \frac{e^{-\bar{\theta}_{0:t}}\left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right) + \alpha e^{\bar{\theta}_{t:T}}\left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right)}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}},$$

$$b_t = \beta \frac{e^{\bar{\theta}_{t:T}} \left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right)}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}},$$

and

$$c_t = 1 - \frac{e^{-\bar{\theta}_{0:t}} \left(\hat{\sigma}^2 + \bar{\sigma}^2_{t:T}\right) + (1-\gamma) e^{\bar{\theta}_{t:T}} \left(\bar{\sigma}^2_{0:T} - \bar{\sigma}^2_{t:T}\right)}{\hat{\sigma}^2 + \bar{\sigma}^2_{0:T}}.$$

We construct a Gaussian one-step forward transition that is consistent with the marginals in Eq. (74). Specifically, we define

$$\begin{aligned} &p_h(\mathbf{x}_t \mid \mathbf{x}_{t-1}) \\ =&\mathcal{N}\left(\mathbf{x}_t; \frac{a_t}{a_{t-1}}\mathbf{x}_{t-1} + \left(b_t - \frac{a_t b_{t-1}}{a_{t-1}}\right)\mathbf{x}^\star + \left(c_t - \frac{a_t c_{t-1}}{a_{t-1}}\right)\boldsymbol{\mu}, \left(\bar{\sigma}'^2_{0:t} - \frac{a_t^2}{a_{t-1}^2}\bar{\sigma}'^2_{0:t-1}\right)\mathbf{I}\right). \end{aligned} \tag{76}$$

This transition is valid whenever

$$\bar{\sigma}'^2_{0:t} - \frac{a_t^2}{a_{t-1}^2}\bar{\sigma}'^2_{0:t-1} \geq 0.$$

It is also consistent with the desired forward marginal. Indeed, if

$$\mathbf{x}_{t-1} \mid \mathbf{x}_0 \sim \mathcal{N}\left(\bar{\boldsymbol{\mu}}'_{0:t-1}, \bar{\sigma}'^2_{0:t-1}\mathbf{I}\right),$$

then Eq. (76) gives

$$\begin{aligned} \mathbb{E}[\mathbf{x}_t \mid \mathbf{x}_0] &= \mathbb{E}\left[\mathbb{E}[\mathbf{x}_t|\mathbf{x}_0, \mathbf{x}_{t-1}]|\mathbf{x}_0\right] \\ &= \mathbb{E}\left[\frac{a_t}{a_{t-1}}\mathbf{x}_{t-1} + \left(b_t - \frac{b_{t-1}a_t}{a_{t-1}}\right)\mathbf{x}^\star + \left(c_t - \frac{a_t c_{t-1}}{a_{t-1}}\right)\boldsymbol{\mu}\middle|\mathbf{x}_0\right] \\ &= \frac{a_t}{a_{t-1}}\bar{\boldsymbol{\mu}}'_{0:t-1} + \left(b_t - \frac{a_t b_{t-1}}{a_{t-1}}\right)\mathbf{x}^\star + \left(c_t - \frac{a_t c_{t-1}}{a_{t-1}}\right)\boldsymbol{\mu} \\ &= a_t\mathbf{x}_0 + b_t\mathbf{x}^\star + c_t\boldsymbol{\mu} = \bar{\boldsymbol{\mu}}'_{0:t}, \end{aligned} \tag{77}$$

and

$$\begin{aligned} \mathrm{Var}(\mathbf{x}_t \mid \mathbf{x}_0) &= \mathbb{E}\left[\mathrm{Var}\left(\mathbf{x}_t|\mathbf{x}_0, \mathbf{x}_{t-1}\right)|\mathbf{x}_0\right] + \mathrm{Var}\left(\mathbb{E}\left[\mathbf{x}_t|\mathbf{x}_0, \mathbf{x}_{t-1}\right]|\mathbf{x}_0\right) \\ &= \left(\bar{\sigma}'^2_{0:t} - \frac{a_t^2}{a_{t-1}^2}\bar{\sigma}'^2_{0:t-1}\right)\mathbf{I} + \mathrm{Var}\left(\frac{a_t}{a_{t-1}}\mathbf{x}_{t-1} + \left(b_t - \frac{a_t b_{t-1}}{a_{t-1}}\right)\mathbf{x}^\star + \left(c_t - \frac{a_t c_{t-1}}{a_{t-1}}\right)\boldsymbol{\mu}\middle|\mathbf{x}_0\right) \\ &= \left(\bar{\sigma}'^2_{0:t} - \frac{a_t^2}{a_{t-1}^2}\bar{\sigma}'^2_{0:t-1}\right)\mathbf{I} + \frac{a_t^2}{a_{t-1}^2}\bar{\sigma}'^2_{0:t-1}\mathbf{I} = \bar{\sigma}'^2_{0:t}\mathbf{I}. \end{aligned} \tag{78}$$

Thus, the transition in Eq. (76) preserves the prescribed forward marginal in Eq. (74).

We now derive the posterior $p_h(\mathbf{x}_{t-1} \mid \mathbf{x}_0, \mathbf{x}_t)$. Under the transition in Eq. (76), we can write

$$\mathbf{x}_t = \frac{a_t}{a_{t-1}}\mathbf{x}_{t-1} + \left(b_t - \frac{a_t b_{t-1}}{a_{t-1}}\right)\mathbf{x}^\star + \left(c_t - \frac{a_t c_{t-1}}{a_{t-1}}\right)\boldsymbol{\mu} + \boldsymbol{\xi}_t,$$

where

$$\boldsymbol{\xi}_t \sim \mathcal{N}\left(\mathbf{0}, \left(\bar{\sigma}'^2_{0:t} - \frac{a_t^2}{a_{t-1}^2}\bar{\sigma}'^2_{0:t-1}\right)\mathbf{I}\right).$$

Moreover, $\boldsymbol{\xi}_t$ is independent of $\mathbf{x}_{t-1}$ conditioned on $\mathbf{x}_0$. In addition, Eq. (74) implies

$$\mathbf{x}_{t-1} \mid \mathbf{x}_0 \sim \mathcal{N}\left(\bar{\boldsymbol{\mu}}'_{0:t-1}, \bar{\sigma}'^2_{0:t-1}\mathbf{I}\right).$$

Hence, $(\mathbf{x}_{t-1}, \mathbf{x}_t) \mid \mathbf{x}_0$ is jointly Gaussian. Its conditional mean is

$$\begin{bmatrix} \bar{\boldsymbol{\mu}}'_{0:t-1} \\ \bar{\boldsymbol{\mu}}'_{0:t} \end{bmatrix}.$$

Moreover, by Eq. (74),

$$\mathrm{Var}(\mathbf{x}_{t-1} \mid \mathbf{x}_0) = \bar{\sigma}'^2_{0:t-1}\mathbf{I}.$$

Hence, by Eq. (76),

$$\begin{aligned}\mathrm{Cov}\left(\mathbf{x}_{t-1}, \mathbf{x}_t \mid \mathbf{x}_0\right) &= \mathrm{Cov}\left(\mathbf{x}_{t-1}, \frac{a_t}{a_{t-1}}\mathbf{x}_{t-1} + \boldsymbol{\xi}_t \middle| \mathbf{x}_0\right) \\ &= \frac{a_t}{a_{t-1}}\mathrm{Var}\left(\mathbf{x}_{t-1} \mid \mathbf{x}_0\right) \\ &= \frac{a_t}{a_{t-1}}\bar{\sigma}'^2_{0:t-1}\mathbf{I}.\end{aligned}$$

Combined with Eq. (74), we obtain that the conditional covariance of $(\mathbf{x}_{t-1}, \mathbf{x}_t) \mid \mathbf{x}_0$ is

$$\begin{bmatrix} \bar{\sigma}'^2_{0:t-1}\mathbf{I} & \frac{a_t}{a_{t-1}}\bar{\sigma}'^2_{0:t-1}\mathbf{I} \\ \frac{a_t}{a_{t-1}}\bar{\sigma}'^2_{0:t-1}\mathbf{I} & \bar{\sigma}'^2_{0:t}\mathbf{I} \end{bmatrix}.$$

Therefore, by the standard conditioning formula for jointly Gaussian random variables,

$$p_h(\mathbf{x}_{t-1} \mid \mathbf{x}_0, \mathbf{x}_t) = \mathcal{N}\left(\mathbf{x}_{t-1}; \boldsymbol{\mu}_{t-1}, \sigma^2_{t-1}\mathbf{I}\right),$$

where

$$\boldsymbol{\mu}_{t-1} = \bar{\boldsymbol{\mu}}'_{0:t-1} + \frac{a_t\bar{\sigma}'^2_{0:t-1}}{a_{t-1}\bar{\sigma}'^2_{0:t}}\left(\mathbf{x}_t - \bar{\boldsymbol{\mu}}'_{0:t}\right),$$

and

$$\sigma^2_{t-1} = \bar{\sigma}'^2_{0:t-1} - \frac{a_t^2\bar{\sigma}'^4_{0:t-1}}{a_{t-1}^2\bar{\sigma}'^2_{0:t}}.$$

Next, we parameterize the reverse transition as

$$p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t) = \mathcal{N}\left(\mathbf{x}_{t-1}; \boldsymbol{\mu}_{\theta,t-1}, \sigma^2_{\theta,t-1}\mathbf{I}\right). \tag{79}$$

Following the Euler discretization of the reverse-time dynamics, and using the noise-prediction parameterization

$$\nabla_{\mathbf{x}_t}\log p_h(\mathbf{x}_t) \approx -\frac{\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, \mathbf{x}^\star, t)}{\bar{\sigma}'_{0:t}}, \tag{80}$$

we set

$$\boldsymbol{\mu}_{\theta,t-1} = \mathbf{x}_t - \left[f_t\mathbf{x}_t + m_t\mathbf{x}^\star + h_t\boldsymbol{\mu} + \frac{\eta_t^2}{\bar{\sigma}'_{0:t}}\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, \mathbf{x}^\star, t)\right], \tag{81}$$

and

$$\sigma^2_{\theta,t-1} = \eta_t^2. \tag{82}$$

Here, $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, \mathbf{x}^\star, t)$ is a neural network parameterized by θ that predicts the standard Gaussian noise component. If a discretization step size Δt is used, then the drift term in Eq. (81) should be multiplied by Δt, and the model variance should be set to $\sigma^2_{\theta,t-1} = \eta_t^2\Delta t$.

Since $\hat{\sigma}^2 > 0$ and $\bar{\sigma}^2_{0:t} > 0$ for any $t > 0$, Eq. (75) implies $\bar{\sigma}'^2_{0:t} > 0$ for all $t \in (0, T]$. In particular, at the terminal time, $\bar{\sigma}'^2_{0:T} > 0$. Therefore, unlike hard bridge parameterizations whose noise scale vanishes at the terminal boundary, our parameterization does not encounter a division-by-zero singularity as t approaches T. This contrasts with the noise-parameterized score used in methods such as GOUB and UniDB, where the endpoint behavior requires special numerical treatment.

Training matches the true posterior $p_h(\mathbf{x}_{t-1} \mid \mathbf{x}_0, \mathbf{x}_t)$ with the parameterized reverse transition $p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t)$. The variational training objective minimizes the KL divergence between these two distributions:

$$\mathrm{KL}\left(p_h(\mathbf{x}_{t-1} \mid \mathbf{x}_0, \mathbf{x}_t) \,\|\, p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t)\right). \tag{83}$$

Using the KL divergence between two isotropic Gaussian distributions, we obtain

$$
\begin{aligned}
&\mathrm{KL}\left(p_h(\mathbf{x}_{t-1}|\mathbf{x}_0,\mathbf{x}_t)||p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)\right)\\
=&\mathbb{E}_{p_h(\mathbf{x}_{t-1}|\mathbf{x}_0,\mathbf{x}_t)}\left[\log\frac{\frac{1}{(2\pi\sigma_{t-1}^2)^{\frac{d}{2}}}\exp\left[-||\mathbf{x}_{t-1}-\boldsymbol{\mu}_{t-1}||^2/2\sigma_{t-1}^2\right]}{\frac{1}{(2\pi\sigma_{\theta,t-1}^2)^{\frac{d}{2}}}\exp\left[-||\mathbf{x}_{t-1}-\boldsymbol{\mu}_{\theta,t-1}||^2/2\sigma_{\theta,t-1}^2\right]}\right]\\
=&\mathbb{E}_{p_h(\mathbf{x}_{t-1}|\mathbf{x}_0,\mathbf{x}_t)}\left[d\log\sigma_{\theta,t-1}-d\log\sigma_{t-1}-\frac{||\mathbf{x}_{t-1}-\boldsymbol{\mu}_{t-1}||^2}{2\sigma_{t-1}^2}+\frac{||\mathbf{x}_{t-1}-\boldsymbol{\mu}_{\theta,t-1}||^2}{2\sigma_{\theta,t-1}^2}\right]\\
=&d\log\sigma_{\theta,t-1}-d\log\sigma_{t-1}-\frac{d}{2}+\frac{d\sigma_{t-1}^2}{2\sigma_{\theta,t-1}^2}+\frac{||\boldsymbol{\mu}_{t-1}-\boldsymbol{\mu}_{\theta,t-1}||^2}{2\sigma_{\theta,t-1}^2}.
\end{aligned}
$$

When $\sigma_{\theta,t-1}^2$ is fixed and not learned, the first four terms on the right-hand side of the above equality are independent of the network parameters. Dropping these constants yields the training objective

$$
\mathcal{L}=\mathbb{E}_{t,\mathbf{x}_0,\mathbf{x}_t}\left[\frac{1}{2\sigma_{\theta,t-1}^2}\left\|\boldsymbol{\mu}_{t-1}-\boldsymbol{\mu}_{\theta,t-1}\right\|^2\right]. \tag{84}
$$

In practice, following prior image-to-image diffusion models (Yue et al., 2024; Zhu et al., 2025; Luo et al., 2023), we omit the coefficient $\frac{1}{2\sigma_{\theta,t-1}^2}$ and replace the squared L_2 mean-matching term with an L_1 surrogate to better preserve pixel-level details. This completes the derivation.

G. GOU Process

Theorem 1. *Consider the following GOU process:*

$$
\mathrm{d}\mathbf{x}_t=\theta_t\left(\boldsymbol{\mu}-\mathbf{x}_t\right)\mathrm{d}t+g_t\mathrm{d}\mathbf{w}_t \tag{85}
$$

where $\boldsymbol{\mu}$ is a given state vector, θ_t denotes a scalar coefficient and g_t represents the diffusion coefficient. We assume the parameterization $g_t^2=2\lambda^2\theta_t$ with $\theta_t>0$, where $\lambda>0$ is a constant. It admits a closed-form analytical solution:

$$
p(\mathbf{x}_t|\mathbf{x}_s)=\mathcal{N}\left(\mathbf{x}_t;\boldsymbol{\mu}+(\mathbf{x}_s-\boldsymbol{\mu})e^{-\bar{\theta}_{s:t}},\frac{g_t^2}{2\theta_t}\left(1-e^{-2\bar{\theta}_{s:t}}\right)\mathbf{I}\right),\ \ \bar{\theta}_{s:t}=\int_s^t\theta_\tau\mathrm{d}\tau \tag{86}
$$

Proof. We define $y_t:=e^{\int_0^t\theta_u\,\mathrm{d}u}$, and $\mathbf{z}_t=\mathbf{x}_t y_t$. Thus, with Itô's formula, we get:

$$
\begin{aligned}
\mathrm{d}\mathbf{z}_t&=\mathrm{d}(\mathbf{x}_t y_t)\\
&=\mathbf{x}_t\mathrm{d}y_t+y_t\mathrm{d}\mathbf{x}_t\\
&=\theta_t\mathbf{x}_t y_t\mathrm{d}t+y_t\left[\theta_t(\boldsymbol{\mu}-\mathbf{x}_t)\mathrm{d}t+g_t\mathrm{d}\mathbf{w}_t\right]\\
&=\boldsymbol{\mu}\theta_t y_t\mathrm{d}t+g_t y_t\mathrm{d}\mathbf{w}_t
\end{aligned} \tag{87}
$$

Integrating from s to t, we get:

$$
\begin{aligned}
\int_s^t\mathrm{d}\mathbf{z}_u&=\boldsymbol{\mu}\int_s^t\theta_u y_u\,\mathrm{d}u+\int_s^t g_u y_u\,\mathrm{d}\mathbf{w}_u\\
\mathbf{z}_t&=\mathbf{z}_s+\boldsymbol{\mu}\int_s^t\theta_u y_u\,\mathrm{d}u+\int_s^t g_u y_u\,\mathrm{d}\mathbf{w}_u\\
\mathbf{x}_t&=\frac{y_s}{y_t}\mathbf{x}_s+\frac{\boldsymbol{\mu}}{y_t}\int_s^t\theta_u y_u\,\mathrm{d}u+\frac{1}{y_t}\int_s^t g_u y_u\,\mathrm{d}\mathbf{w}_u\\
&=e^{-\int_s^t\theta_u\,\mathrm{d}u}\mathbf{x}_s+\boldsymbol{\mu}\int_s^t e^{-\int_u^t\theta_\tau\,\mathrm{d}\tau}\theta_u\,\mathrm{d}u+\int_s^t e^{-\int_u^t\theta_\tau\,\mathrm{d}\tau}g_u\,\mathrm{d}\mathbf{w}_u\\
&=e^{-\int_s^t\theta_u\,\mathrm{d}u}\mathbf{x}_s+\boldsymbol{\mu}\left[1-e^{-\int_s^t\theta_u\,\mathrm{d}u}\right]+\int_s^t e^{-\int_u^t\theta_\tau\,\mathrm{d}\tau}g_u\,\mathrm{d}\mathbf{w}_u
\end{aligned} \tag{88}
$$

It is obvious that the transition kernel is a Gaussian distribution. And we have:

$$\begin{aligned}\int_s^t e^{-\int_u^t \theta_\tau \,\mathrm{d}\tau} g_u \,\mathrm{d}\mathbf{w}_u &\sim \mathcal{N}\left(\mathbf{0}, \int_s^t e^{-2\int_u^t \theta_\tau \,\mathrm{d}\tau} g_u^2 \,\mathrm{d}u\, \mathbf{I}\right) \\ &= \mathcal{N}\left(\mathbf{0}, \lambda^2 \int_s^t e^{-2\int_u^t \theta_\tau \,\mathrm{d}\tau} 2\theta_u \,\mathrm{d}u\, \mathbf{I}\right) \\ &= \mathcal{N}\left(\mathbf{0}, \frac{g_t^2}{2\theta_t}\left(1 - e^{-2\bar{\theta}_{s:t}}\right)\mathbf{I}\right)\end{aligned} \tag{89}$$

Therefore:

$$p(\mathbf{x}_t \mid \mathbf{x}_s) = \mathcal{N}\left(\mathbf{x}_t; e^{-\bar{\theta}_{s:t}}\mathbf{x}_s + \left(1 - e^{-\bar{\theta}_{s:t}}\right)\boldsymbol{\mu}, \frac{g_t^2}{2\theta_t}\left(1 - e^{-2\bar{\theta}_{s:t}}\right)\mathbf{I}\right). \tag{90}$$

This concludes the proof of Theorem 1.

□

H. Experimental Details

We follow the network architecture and evaluation protocols of GOUB (Yue et al., 2024) and UniDB (Zhu et al., 2025), with the training configuration specified below. We parameterize the noise estimation network using a standard U-Net architecture (Luo et al., 2023). The models are optimized using Adam (Kingma & Ba, 2015) with momentum parameters $\beta_1 = 0.9$ and $\beta_2 = 0.99$, and a learning rate of 10^{-4}. During training, we use randomly cropped 128×128 image patches with a batch size of 8. For the noise schedule θ_t, we adopt the flipped cosine schedule (Nichol & Dhariwal, 2021; Luo et al., 2023):

$$\theta_t = 1 - \frac{\cos^2\left(\frac{t/T+s}{1+s}\frac{\pi}{2}\right)}{\cos^2\left(\frac{s}{1+s}\frac{\pi}{2}\right)}, \tag{91}$$

where the offset $s = 0.008$ is applied to ensure a smooth transition near the boundary. Following established protocols, we set $\lambda = 30/255$ by default. For larger σ values in the ablation study, we increase λ as needed to satisfy $\sigma^2 < \bar{\sigma}_{0:T}^2$. The best-performing σ values lie within the admissible range under the default λ. To maintain numerical stability, the terminal coefficient $e^{-\bar{\theta}_{0:T}}$ is truncated at 0.005 rather than strictly zero.

All models are trained on a single NVIDIA A100 GPU (40GB). For the deraining and inpainting tasks, we train the models for 1.2 million iterations. For the super-resolution task, the training budget is set to $600,000$ iterations.

During sampling, we initialize the reverse process with $\mathbf{x}_T \approx (b_T + c_T)\mathbf{x}^\star$, omitting the unavailable source-dependent term and the terminal noise.

For sampling, we use the following reverse process,

$$\mathrm{d}\mathbf{x}_t = \left[\mathbf{f}(\mathbf{x}_t, t) + \zeta\eta_t^2 \frac{\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, \mathbf{x}^\star, t)}{\bar{\sigma}'_{0:t}}\right]\mathrm{d}t + \eta_t \,\mathrm{d}\bar{\mathbf{w}}_t,\ \mathrm{d}\mathbf{x}_t = \left[\mathbf{f}(\mathbf{x}_t, t) + \frac{1}{2}\zeta\eta_t^2 \frac{\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, \mathbf{x}^\star, t)}{\bar{\sigma}'_{0:t}}\right]\mathrm{d}t \tag{92}$$

When $\zeta = 1$, the above formulations recover the standard reverse SDE and its corresponding probability flow ODE. However, Eq. (81) is derived from a first-order Euler discretization of the reverse SDE, which introduces discretization errors. Therefore, directly applying the theoretically derived drift coefficient may lead to inaccurate score estimation and unstable sampling trajectories in practice. To compensate for this discrepancy, we introduce ζ as an empirical correction factor for the score-driven drift term. As shown in Table 3, adjusting ζ provides improved restoration performance and more stable sampling trajectories compared with using the theoretical coefficient. Therefore, we treat ζ as an empirical coefficient and set it to 1.12, 1.115, and 1.045 for deraining, inpainting, and super-resolution, respectively.

Table 3: Quantitative comparison of different sampling strategies on image restoration tasks. Results are reported on DIV2K, Rain100H, and CelebA-HQ 256×256 datasets for image super-resolution, deraining, and inpainting, respectively.

METHOD	Image Deraining				METHOD	Image Super-Resolution				METHOD	Image Inpainting			
	PSNR↑	SSIM↑	LPIPS↓	FID↓		PSNR↑	SSIM↑	LPIPS↓	FID↓		PSNR↑	SSIM↑	LPIPS↓	FID↓
SDE	32.38	0.9069	0.042	15.68	SDE	22.81	0.5409	0.366	33.78	SDE	29.62	0.9084	0.034	3.86
Mean-ODE	34.96	0.9450	0.069	28.44	Mean-ODE	28.80	0.8141	0.314	20.30	Mean-ODE	32.00	**0.9404**	0.051	11.84
SDE-ζ	32.94	0.9183	**0.038**	**14.21**	SDE-ζ	27.22	0.7522	**0.136**	**13.70**	SDE-ζ	30.30	0.9201	**0.032**	**3.48**
ODE-ζ	**35.06**	**0.9464**	0.066	25.75	ODE-ζ	**28.90**	**0.8164**	0.309	19.95	ODE-ζ	**32.22**	0.9399	0.051	11.87

I. Ablation Study

In this section, we systematically evaluate the sensitivity of our framework to its core hyperparameters. Recall that the h-transform parameters α, β, and γ modulate the soft terminal center location, while $\hat{\sigma}^2$ determines its variance spread. We omit γ from this analysis because we consistently set $\boldsymbol{\mu} = \mathbf{x}^\star$ across all experiments, rendering its independent variation redundant. Furthermore, since $\hat{\sigma}^2$ is explicitly determined by σ via Eq. (12), our ablation study focuses exclusively on the three essential parameters: α, β, and σ.

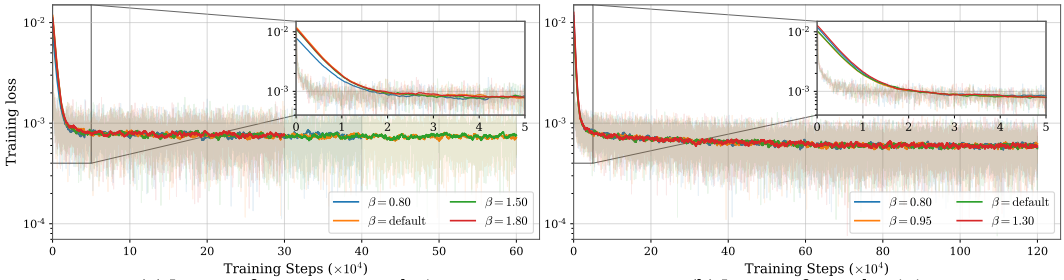

(a) Impact of β on super-resolution (b) Impact of β on deraining

Figure 5: Training loss for different values of β on the super-resolution and deraining tasks.

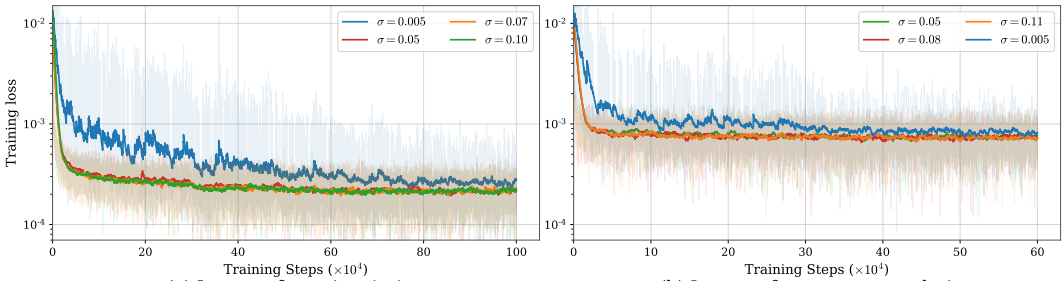

(a) Impact of σ on inpainting (b) Impact of σ on super-resolution

Figure 6: Training loss for different values of σ on the inpainting and super-resolution tasks.

Sensitivity Analysis of σ: We first study the effect of the terminal variance parameter σ^2, which controls the strength of the prescribed Gaussian terminal constraint in SDDBMs. As shown in Figure 2(a), the performance first improves as σ increases, but degrades when σ becomes too large. An overly small σ makes the terminal distribution highly concentrated, causing the process to approach a hard endpoint bridge and become more susceptible to terminal sampling instability. In contrast, an excessively large σ over-relaxes the terminal constraint and weakens target-guided transport. Thus, σ governs a key trade-off between sampling stability and target fidelity. Figure 4 further supports this observation from a training perspective, showing that both extremely small and very large values of σ lead to higher training losses, while moderate values stabilize training. These results in Figure 4 suggest that $\sigma \in [0.05, 0.11]$ is a reasonably stable range, while the best performance is typically achieved for $\sigma \in [0.08, 0.11]$. Similar trends are also observed on the other two tasks, as shown in Figure 6(a)(b).

Impact of α and β: The quantitative results regarding hyperparameter α are depicted in Figure 2(b). Notably, the restoration performance degrades precipitously as α approaches the critical threshold of $-e^{-\bar{\theta}_{0:T}}\hat{\sigma}^2/\bar{\sigma}^2_{0:T}$, as established in Lemma 2. This severe drop empirically validates that nearing this limit triggers the singularity problem, substantially undermining the model's generative efficacy. Regarding β (Figure 5), we observe that different β values lead to comparable performance, indicating that the model is robust to the choice of β within the tested range. However, a relatively large β tends to produce higher training losses during the early stage of optimization and slightly degrades the final performance. Therefore, we recommend choosing β values around the default setting[2], which provides stable training and reliable restoration quality.

J. Additional Visual Results

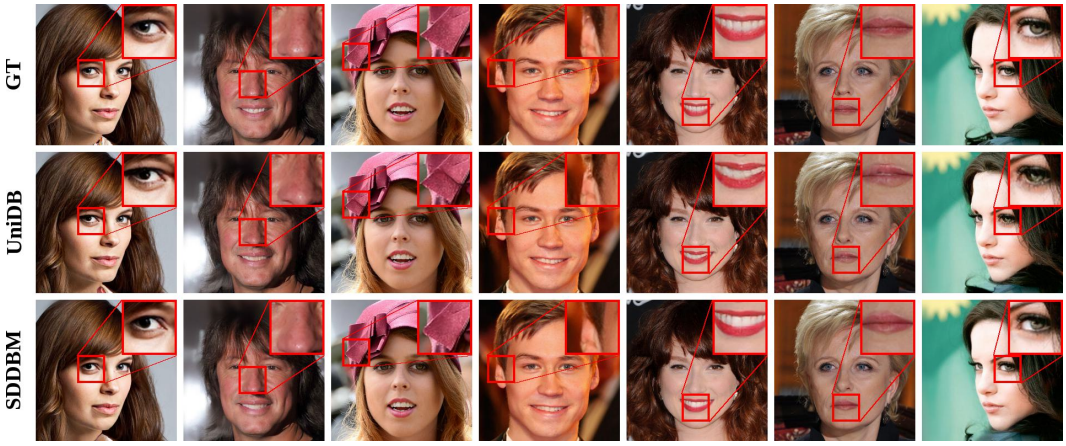

Figure 7: Qualitative comparison between UniDB and SDDBMs on the CelebA-HQ dataset for image inpainting.

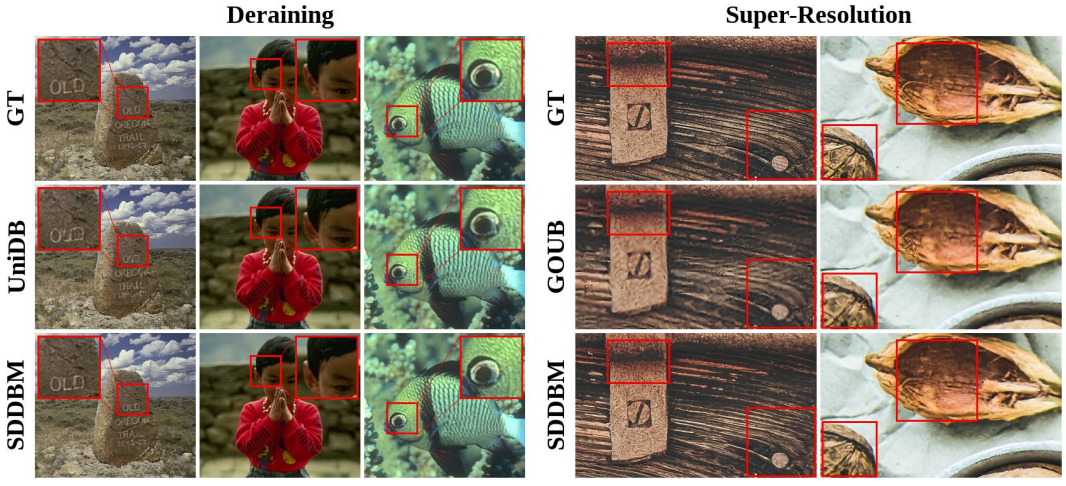

Figure 8: Qualitative comparison between SDDBMs and baseline methods on the Rain100H dataset for image deraining (left) and the DIV2K dataset for image super-resolution (right).

[2]By default, we set $\beta = 1 + \frac{\hat{\sigma}^2}{\bar{\sigma}^2_{0:T}}$ and $\gamma = -\frac{\hat{\sigma}^2}{\bar{\sigma}^2_{0:T}}\left(1 - e^{-\bar{\theta}_{0:T}}\right)$, yielding $b_T = 1$.

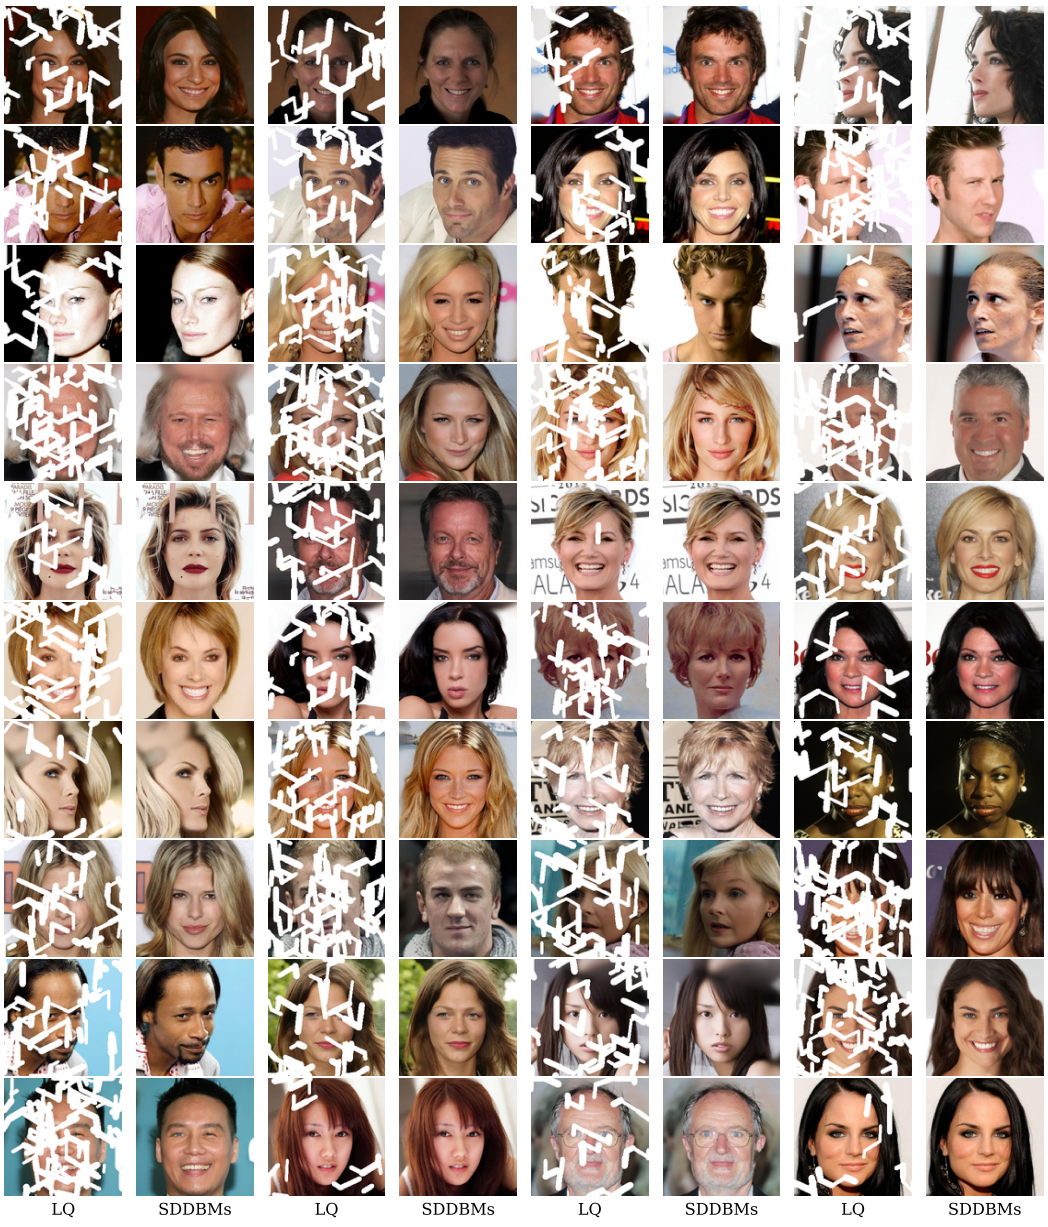

Figure 9: Additional visual results for thin-mask inpainting on the CelebA-HQ dataset.



Figure 10: Additional visual results for deraining on the Rain100H dataset.

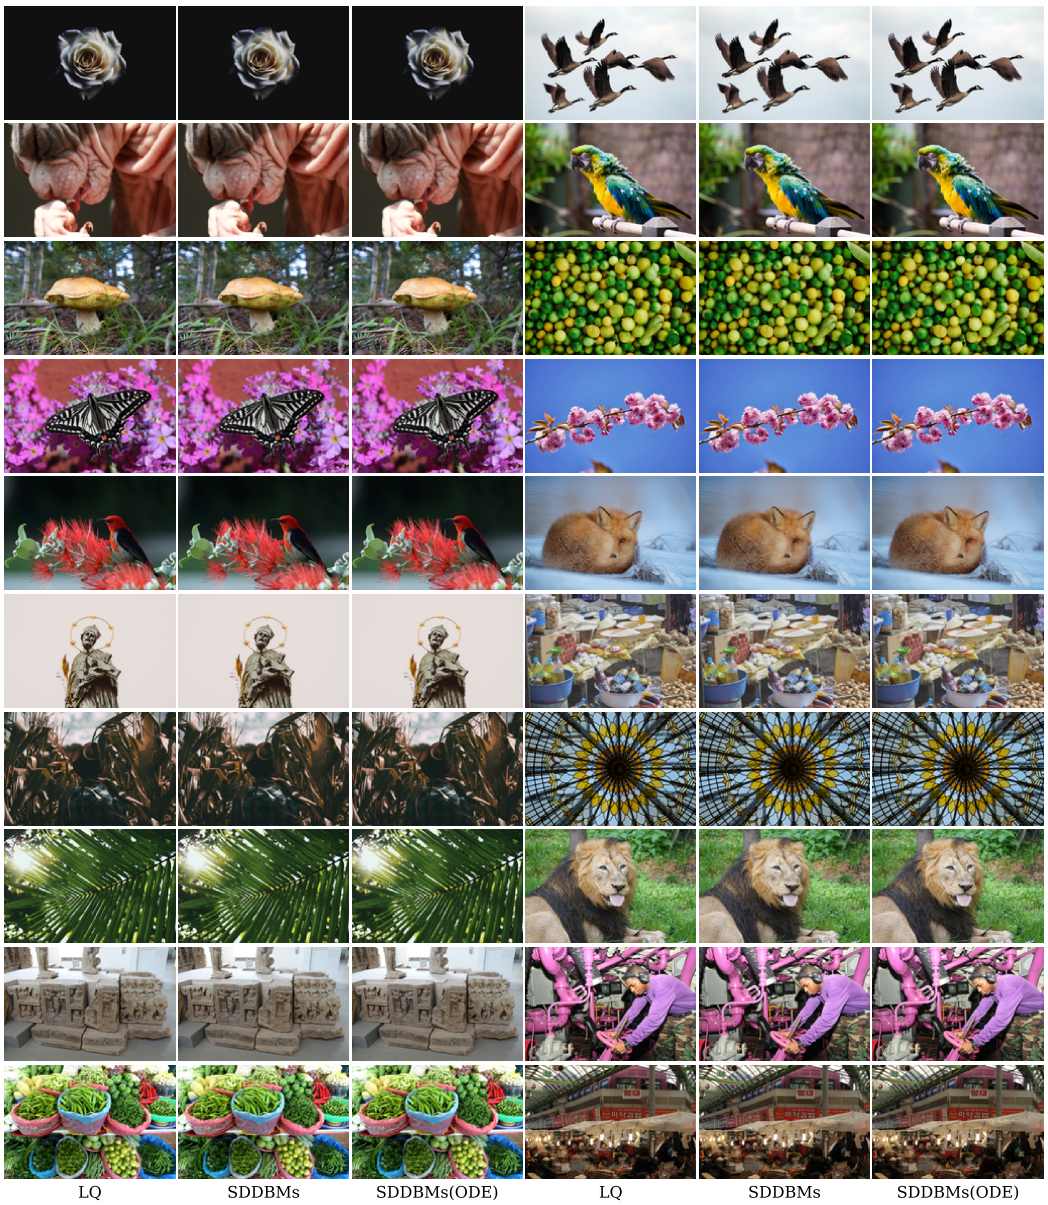

Figure 11: Additional visual results for super-resolution on the DIV2K dataset.