Deterministic Discrete Denoising

Abstract

We propose a deterministic denoising algorithm for discrete-state diffusion models. The key idea is to derandomize the generative reverse Markov chain by introducing a variant of the herding algorithm, which induces deterministic state transitions driven by weakly chaotic dynamics. It serves as a direct replacement for the stochastic denoising process, without requiring retraining or continuous state embeddings. We demonstrate consistent improvements in both efficiency and sample quality on text and image generation tasks. In addition, the proposed algorithm yields improved solutions for diffusion-based combinatorial optimization. Thus, herding-based denoising is a simple yet promising approach for enhancing the generative process of discrete diffusion models. Furthermore, our results reveal that deterministic reverse processes, well established in continuous diffusion, can also be effective in discrete state spaces.

1. Introduction

Diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020) have recently achieved remarkable success in generating realistic image and audio data. Their forward process progressively corrupts data with noise, and the models are trained to denoise by inverting the corruption process. Once trained, diffusion models can generate realistic samples from pure noise by iteratively applying the denoising procedure. For continuous data such as images and speech, the reverse denoising process is often implemented as a deterministic algorithm, which exhibits more directed dynamics and enables efficient generation with fewer denoising steps compared to stochastic denoising, as demonstrated by DDIM (Song et al., 2021).

Diffusion models for discrete data (Hoogeboom et al., 2021;

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

Austin et al., 2021) are also an important research direction, since many data of interest are inherently discrete, including text, graphs, genomes, and chemical structures. In discrete-state diffusion models, the forward and reverse processes are typically formulated as probabilistic Markov chains on a discrete sample space (Austin et al., 2021). If the reverse denoising process could be made deterministic, improved performance would be expected, as in the case of continuous models. However, in discrete spaces, naive derandomization as a deterministic mapping aggregates discrete trajectories and reduces sample diversity, which is essential for generative modeling.

One promising approach is to embed the discrete sample space into a continuous space and apply continuous diffusion models (Hoogeboom et al., 2021; Chen et al., 2023; Sahoo et al., 2025). The main advantage of this approach is that it allows the direct use of well-established techniques for continuous diffusion, including deterministic denoising. Nevertheless, denoising in the continuous space requires retraining the model. Moreover, it remains unclear whether embedding discrete data into a continuous space is truly necessary for derandomization.

In this paper, we propose to directly derandomize the reverse denoising process on discrete sample spaces. The key idea is to introduce a variant of the herding algorithm (Welling, 2009; Welling & Chen, 2010), which associates auxiliary continuous weight variables with each discrete variable in the sample space. The reverse denoising updates of both discrete and continuous variables are deterministic, and randomness arises solely from their initial states. Since the proposed algorithm derandomizes the state transitions of Markov chains, it serves as a drop-in replacement for the stochastic reverse process without retraining discrete diffusion models. This relation is analogous to DDIM (Song et al., 2021), which generates samples deterministically using a model trained for DDPM (Ho et al., 2020), as summarized in Table 1. In our experiments on text and image generation, we demonstrate consistent improvements in both efficiency and sample quality. These results reveal that deterministic reverse processes, well established in continuous diffusion, can also be effective in discrete state spaces.

For evaluation, we focus on discrete uniform diffusion models, while the proposed denoising algorithm is applicable to

Table 1. Deterministic denoising algorithms for diffusion models with continuous, masked discrete, and uniform discrete noising processes. We establish a herding-based deterministic denoising algorithm for discrete diffusion models without masking.

	Stochastic	Deterministic
Continuous	DDPM (Ho et al., 2020)	DDIM (Song et al., 2021)
Discrete Masked	D3PM absorbing (Austin et al., 2021) MDLM (Sahoo et al., 2024)	DNDM (Chen et al., 2024)
Discrete Uniform	D3PM uniform (Austin et al., 2021) UDLM (Schiff et al., 2025)	**herding-based denoising (ours)**

various noising processes. Specifically, in our experiments, we adopt UDLM (Schiff et al., 2025), a state-of-the-art discrete-state generative model based on the uniform diffusion process. Our aim in this paper is not to propose a new discrete diffusion model, but to demonstrate how the herding-based algorithm improves the generative process of existing models. We use the official training, sampling, and evaluation codes of UDLM, modifying only about 20 lines in the reverse denoising procedure. This directly demonstrates that our method works as a drop-in replacement without retraining. Moreover, the uniform diffusion process is particularly suitable for our approach, since its reverse process performs iterative refinement through successive Markov state transitions, making it a natural setting to examine whether derandomizing these transitions can improve efficiency and sample quality. As another evaluation axis, we apply the proposed algorithm to DIFUSCO (Sun & Yang, 2023), a diffusion-based combinatorial optimization framework, to assess improvements in objective values achieved by the derandomized reverse process.

Diffusion language models are an important research area within discrete diffusion, while our approach is not restricted to text data. In this context, masked diffusion models have been extensively studied and are expected to enable parallel and coherent text generation with iterative refinement. While autoregressive (AR) models (Bengio et al., 2003) generate tokens sequentially, masked diffusion models allow tokens to be generated in arbitrary orders, which can offer flexibility in generation. Early masked diffusion models can be regarded as relatively static in their reverse processes, in the sense that each token transitions only once from the masked state and remains fixed thereafter, which is a consequence of the absorbing noising process toward the all-mask state. As summarized in Table 1, derandomization for such static masked diffusion has been explored in DNDM (Chen et al., 2024), where the reverse process is derandomized by externally specifying the unmasking transition times. More recent studies, however, have moved toward more dynamic

reverse processes that enable iterative refinement, where tokens may be remasked, as in LLaDA (Nie et al., 2025; Zhu et al., 2025), or resampled multiple times, as in GIDD (von Rütte et al., 2025). The uniform diffusion process can be considered the extreme case of this trend, where Markov state transitions in the token space continue throughout the reverse process (Schiff et al., 2025; von Rütte et al., 2025). Notably, recent work suggests that uniform diffusion scales more efficiently than masked diffusion in data-limited large-parameter regimes (von Rütte et al., 2025), which is particularly relevant given the growing concern over data scarcity (Villalobos et al., 2024). Accordingly, in diffusion-based text generation, our derandomization approach is expected to facilitate effective iterative refinement in dynamic reverse processes driven by Markov chains.

Overall, we present a simple yet effective derandomization approach that directly enhances the generative process of existing discrete-state diffusion models. The proposed algorithm exhibits weakly chaotic, or piecewise isometric dynamics, which preserves distance and probability mass locally, with a superior asymptotic convergence rate compared with stochastic sampling. Thus, this work provides a new perspective for viewing discrete diffusion models through the notions of nonlinear dynamics and deterministic sampling.

2. Preliminaries

2.1. Discrete Diffusion Models

A diffusion process in a discrete state space can be viewed as a sequence of discrete state transitions that gradually corrupt the original data. This process can be formulated as a discrete-time Markov chain over categorical variables, giving rise to a reverse-time denoising Markov chain (Austin et al., 2021). This formulation can be naturally extended to state transitions in continuous time, using Markov jump processes parameterized by transition-rate matrices (Campbell et al., 2022; Sun et al., 2023). Foundational theoretical studies analyze approximation error and convergence properties of both discrete-time and continuous-time discrete diffusion models (Ren et al., 2025a; Zhang et al., 2025; Liang et al., 2025a;b). Below, we focus on a discrete-time formulation to outline how our approach can be applied to the reverse denoising process.

We consider generative models for discrete data represented as a sequence of L categorical variables. Each variable takes one of K categories and is represented as a K-dimensional one-hot vector $\boldsymbol{x} \in \mathcal{V}$, where $\mathcal{V} = \{(x_1, x_2, \ldots, x_K)^\top \in \{0,1\}^K \mid \sum_{k=1}^K x_k = 1\}$. A sequence of L variables is denoted by $\boldsymbol{x}^{(1:L)} = (\boldsymbol{x}^{(1)}, \boldsymbol{x}^{(2)}, \ldots, \boldsymbol{x}^{(L)}) \in \mathcal{V}^L$.

The forward noising process in the D3PM framework (Austin et al., 2021) is formulated as a discrete-time Markov

Algorithm 1 Stochastic denoising for uniform diffusion

Require: trained prediction model p_θ
1: **for** $l = 1$ to L **do**
2:   $\boldsymbol{x}_T^{(l)} \sim \text{Uniform}(\mathcal{V})$
3: **end for**
4: **for** $t = T$ down to 1 **do**
5:   **for** $l = 1$ to L **do**
6:     $\boldsymbol{p}_{t-1}^{(l)} \leftarrow p_\theta(\boldsymbol{x}_{t-1}^{(l)} \mid \boldsymbol{x}_t^{(1:L)}, t)$
7:     $\boldsymbol{x}_{t-1}^{(l)} \sim \text{Cat}(\boldsymbol{x}_{t-1}^{(l)}; \boldsymbol{p}_{t-1}^{(l)})$
8:   **end for**
9: **end for**
10: **Output:** generated sample $\boldsymbol{x}_0^{(1:L)}$

---

chain on $\mathcal{V}$, which gradually transforms the original data $\boldsymbol{x}_0$ into noise $\boldsymbol{x}_T$. Given the Markov transition matrix Q_t at time t, the forward transition is defined as

$$q(\boldsymbol{x}_t \mid \boldsymbol{x}_{t-1}) = \text{Cat}(\boldsymbol{x}_t; Q_t\boldsymbol{x}_{t-1}), \qquad (1)$$

where $\text{Cat}(\cdot; \boldsymbol{p})$ denotes the categorical distribution with the probability vector $\boldsymbol{p}$.

The Markov transition matrix Q_t is designed to gradually transform any categorical distribution into a stationary distribution as t increases. For the uniform noising process, it is given by

$$Q_t = (1 - \beta_t)\boldsymbol{I} + \beta_t \frac{1}{K}\boldsymbol{1}\boldsymbol{1}^\top, \qquad (2)$$

where $\beta_t \in (0, 1)$ is a noise schedule. For the absorbing noising process, it is defined as

$$Q_t = (1 - \beta_t)\boldsymbol{I} + \beta_t \boldsymbol{x}_{\text{mask}}\boldsymbol{1}^\top, \qquad (3)$$

where $\boldsymbol{x}_{\text{mask}}$ is the one-hot vector representing the absorbing category. In both cases, the Markov transition matrix from time s to t, $Q_t Q_{t-1} \dots Q_{s+1}$, admits a closed-form expression for arbitrary $s < t$, which can be extended to non-integer times. This property is exploited to improve the efficiency of training and sampling (Sahoo et al., 2024; Shi et al., 2024; Schiff et al., 2025).

The reverse denoising process is also formulated as a discrete-time Markov chain on $\mathcal{V}^L$. A neural network model parameterized by θ is trained to predict $\tilde{\boldsymbol{x}}_0^{(1:L)}$, an estimate of the original data $\boldsymbol{x}_0^{(1:L)}$, from noisy data $\boldsymbol{x}_t^{(1:L)}$ at time t. Training is performed by minimizing a variational upper bound (NELBO) on the negative log-likelihood (Austin et al., 2021; Schiff et al., 2025). The transition probability $p_\theta(\boldsymbol{x}_{t-1} \mid \boldsymbol{x}_t^{(1:L)}, t)$ of the reverse process is then computed from the predicted original data $\tilde{\boldsymbol{x}}_0^{(1:L)}$ and the forward transition matrix Q_t.

By iterating this Markov transition from pure noise $\boldsymbol{x}_T^{(1:L)}$ sampled from the stationary distribution, the reverse process generates a sample $\boldsymbol{x}_0^{(1:L)}$, as described in Algorithm 1.

Namely, each denoising step reduces to sampling from the categorical distribution $\text{Cat}(\boldsymbol{x}_{t-1}; \boldsymbol{p}_{t-1})$ with probability vector $\boldsymbol{p}_{t-1} = p_\theta(\boldsymbol{x}_{t-1} \mid \boldsymbol{x}_t^{(1:L)}, t)$. This is typically implemented as a probabilistic procedure using the Gumbel-max trick. However, sampling does not necessarily have to be probabilistic; it suffices that samples are generated faithfully according to the given transition probabilities. This motivates us to explore deterministic alternatives for the reverse denoising process.

2.2. Herding Algorithm

The herding algorithm (Welling, 2009; Welling & Chen, 2010) is a deterministic sampling method with weakly chaotic dynamics, which combines statistical learning and inference for energy-based models into a single dynamical system.

The herding system yields a sample sequence $\boldsymbol{x}_1, \boldsymbol{x}_2, \dots$ from a discrete sample space $\mathcal{V}$ such that the empirical average of feature values $\boldsymbol{\phi}(\boldsymbol{x}_t)$ converges to predefined target values $\boldsymbol{\mu} \in \mathbb{R}^N$, where $\boldsymbol{\phi}: \mathcal{V} \to \mathbb{R}^N$ is a set of real-valued feature functions.

To reduce the discrepancy between the averaged feature values and the target values in each sampling step, the herding system updates the weight vector $\boldsymbol{w}_t$. Specifically, the herding system generates a new sample $\boldsymbol{x}_{t+1}$ and updates the weight vector $\boldsymbol{w}_{t+1}$ at time $t + 1$ as follows:

$$\boldsymbol{x}_{t+1} = \arg\max_{\boldsymbol{x} \in \mathcal{V}} \boldsymbol{w}_t^\top \boldsymbol{\phi}(\boldsymbol{x}), \qquad (4)$$

$$\boldsymbol{w}_{t+1} = \boldsymbol{w}_t + \boldsymbol{\mu} - \boldsymbol{\phi}(\boldsymbol{x}_{t+1}). \qquad (5)$$

This forms a hybrid dynamical system (Aihara & Suzuki, 2010) consisting of a discrete state variable $\boldsymbol{x}_t$ and a continuous state variable $\boldsymbol{w}_t$. Given an initial state $(\boldsymbol{x}_0, \boldsymbol{w}_0) \in \mathcal{V} \times \mathbb{R}^N$, the system generates samples iteratively, deterministically, and autonomously. Typically, the weight vector $\boldsymbol{w}_t$ does not converge and remains in a bounded domain around the origin.

The difference between the predefined target vector and the feature vector averaged over T samples can be written as

$$\boldsymbol{\mu} - \frac{1}{T}\sum_{t=1}^{T}\boldsymbol{\phi}(\boldsymbol{x}_t) = \frac{1}{T}(\boldsymbol{w}_T - \boldsymbol{w}_0). \qquad (6)$$

The right-hand side converges to zero as $T \to \infty$, provided that the sequence $\boldsymbol{w}_t$ remains bounded, which is guaranteed under mild conditions. Therefore, the herding system generates a sequence of samples with the predefined feature expectations converging asymptotically at a rate of $O(T^{-1})$. This convergence rate is much faster than the $O(T^{-1/2})$ rate of probabilistic sampling, such as independent categorical sampling and Markov chain Monte Carlo (MCMC) sampling, from a distribution with the same feature expectations.

Note that exact maximization in the sample generation step is not required, since the $O(T^{-1})$ convergence holds as long as the weight vector remains bounded (Welling, 2009; Welling & Chen, 2010). The convergence rate of the herding algorithm is analyzed in more detail (Bach et al., 2012; Harvey & Samadi, 2014).

The dynamics of the herding system is classified as a piecewise isometry (Goetz, 2000; 2003), which is weakly chaotic, with Lyapunov exponents equal to zero almost everywhere, and typically has a fractal attracting set. Such complex nonlinear dynamics of the herding system yields diverse samples with relatively high entropy and negative autocorrelations (Welling, 2009; Welling & Chen, 2010).

For sampling from a categorical distribution $\mathrm{Cat}(\boldsymbol{x}; \boldsymbol{p})$ on $\mathcal{V}$, the feature function $\boldsymbol{\phi}(\boldsymbol{x}) = \boldsymbol{x}$ and the target vector $\boldsymbol{\mu} = \boldsymbol{p}$ lead to the following herding dynamics:

$$\boldsymbol{x}_{t+1} = \arg\max_{\boldsymbol{x} \in \mathcal{V}} \boldsymbol{w}_t^\top \boldsymbol{x}, \tag{7}$$

$$\boldsymbol{w}_{t+1} = \boldsymbol{w}_t + \boldsymbol{p} - \boldsymbol{x}_{t+1}. \tag{8}$$

This system generates samples from $\mathrm{Cat}(\boldsymbol{x}; \boldsymbol{p})$ whose empirical distribution is closer to $\boldsymbol{p}$ than that of probabilistic categorical sampling.

3. Herding-Based Denoising

3.1. Derandomization of Reverse Markov Chain

We apply the herding algorithm to derandomize the reverse Markov chain of discrete-state diffusion models. Since the transition probability at each denoising step is time-dependent, we need to introduce a time-dependent version of the herding algorithm as in (Suzuki, 2014).

Specifically, the sample $\boldsymbol{x}_t$ and weight vector $\boldsymbol{w}_t$ are updated from time t to $t - 1$ as follows:

$$\boldsymbol{x}_{t-1} = \arg\max_{\boldsymbol{x} \in \mathcal{V}} (\boldsymbol{w}_t + \boldsymbol{p}_{t-1})^\top \boldsymbol{x}, \tag{9}$$

$$\boldsymbol{w}_{t-1} = \boldsymbol{w}_t + \boldsymbol{p}_{t-1} - \boldsymbol{x}_{t-1}, \tag{10}$$

where $\boldsymbol{p}_{t-1} = p_\theta(\boldsymbol{x}_{t-1} \mid \boldsymbol{x}_t^{(1:L)}, t)$ denotes the probability vector of the categorical distribution obtained from the model prediction. Since the transition probability is time-dependent, the current probability vector $\boldsymbol{p}_{t-1}$ is taken into account when selecting the new sample $\boldsymbol{x}_{t-1}$ in Eq. (9), as opposed to the original herding update in Eq. (7).

Additionally, we introduce a delayed-switching mechanism with a margin $\delta > 0$ to mitigate excessive switching of the discrete state variable, by modifying Eq. (9) as follows:

$$\boldsymbol{x}_{t-1} = \arg\max_{\boldsymbol{x} \in \mathcal{V}} (\boldsymbol{w}_t + \boldsymbol{p}_{t-1} + \delta \boldsymbol{x}_t)^\top \boldsymbol{x}. \tag{11}$$

That is, the sample $\boldsymbol{x}_{t-1}$ remains the same as $\boldsymbol{x}_t$ unless

---

Algorithm 2 Herding-based denoising

Require: trained prediction model p_θ
1: **for** $l = 1$ to L **do**
2:     $\boldsymbol{x}_T^{(l)} \sim \mathrm{Uniform}(\mathcal{V})$
3:     $\boldsymbol{w}_T^{(l)} \sim \mathrm{Uniform}([0, 1]^K)$
4: **end for**
5: **for** $t = T$ down to 1 **do**
6:     **for** $l = 1$ to L **do**
7:        $\boldsymbol{p}_{t-1}^{(l)} \leftarrow p_\theta(\boldsymbol{x}_{t-1}^{(l)} \mid \boldsymbol{x}_t^{(1:L)}, t)$
8:        $\boldsymbol{x}_{t-1}^{(l)} \leftarrow \arg\max_{\boldsymbol{x} \in \mathcal{V}} (\boldsymbol{w}_t^{(l)} + \boldsymbol{p}_{t-1}^{(l)} + \delta \boldsymbol{x}_t^{(l)})^\top \boldsymbol{x}$
9:        $\boldsymbol{w}_{t-1}^{(l)} \leftarrow \boldsymbol{w}_t^{(l)} + \boldsymbol{p}_{t-1}^{(l)} - \boldsymbol{x}_{t-1}^{(l)}$
10:     **end for**
11: **end for**
12: **Output:** generated sample $\boldsymbol{x}_0^{(1:L)}$

---

another candidate exceeds the objective value by at least δ. A pseudocode of this algorithm is provided in Algorithm 2.

This herding-based system reduces the cumulative discrepancy $\sum_{\tau=t}^{T-1} (\boldsymbol{p}_\tau - \boldsymbol{x}_\tau)$ compared with stochastic sampling, thereby yielding a sample sequence that reflects transition probabilities more faithfully. The boundedness and $O(T^{-1})$ discrepancy bound, inherited from the original herding system, are proved in Section 3.2.

As the proposed denoising algorithm is entirely deterministic, the only source of randomness arises from the initialization of $(\boldsymbol{x}_T, \boldsymbol{w}_T)$, which is drawn randomly from the uniform distribution on $\mathcal{V} \times [0, 1]^K$. Through the iterative updates, the algorithm induces a mapping from a pure noise $(\boldsymbol{x}_T^{(1:L)}, \boldsymbol{w}_T^{(1:L)})$ to a sample $(\boldsymbol{x}_0^{(1:L)}, \boldsymbol{w}_0^{(1:L)})$, which constitutes a piecewise isometry. Unlike the mapping defined by ODE flows in continuous diffusion models, this mapping is not one-to-one, and the probability mass is preserved piecewise. In other words, small perturbations in the weight vector persist unless they alter the argmax. Nevertheless, once the generated sample changes, these perturbations cause a drastic shift in the trajectory, reflecting the weakly chaotic dynamics.

The time-dependent herding system introduced here converts a K-dimensional real-valued sequence $\{\boldsymbol{p}_t\}$ into a one-hot discrete sequence $\{\boldsymbol{x}_t\}$. In this sense, it can also be viewed as a multidimensional extension of 1-bit analog-to-digital (A/D) converters that represent continuous signals using discrete binary signals, as exemplified by the $\Delta\Sigma$-modulator (Inose et al., 1962; Inose & Yasuda, 1963) and the error diffusion algorithm (Adler et al., 2003), which have long been studied in engineering fields such as signal and image processing.

Figure 1. Evaluation of text8 character-level text samples by generative PPL (Gen. PPL) and entropy. Results for stochastic denoising (Gumbel-max) and herding-based denoising with different δ values are shown.

3.2. Bounded Dynamics, Convergence, and Initialization

As a theoretical justification for our proposed algorithm, we provide a proof of the boundedness of the weight vector $\boldsymbol{w}_t$ and the $O(T^{-1})$ bound on the discrepancy between the empirical sample distribution and the averaged probability vector.

Proposition 1. *For any initial weight $\boldsymbol{w}_T \in \mathbb{R}^K$, the infinite reverse orbit $\{\boldsymbol{w}_t\}_{t \leq T}$ generated by the herding-based denoising is bounded.*

Proof. Let $w_{t,i}, p_{t,i}, x_{t,i}$ denote the ith component of $\boldsymbol{w}_t$, $\boldsymbol{p}_t$, $\boldsymbol{x}_t$, respectively. We first observe that the total weight $\sum_i w_{t,i}$ is conserved, as

$$\sum_i w_{t-1,i} = \sum_i (w_{t,i} + p_{t-1,i} - x_{t-1,i})$$
$$= \sum_i w_{t,i} + 1 - 1 = \sum_i w_{t,i}.$$

Let $\mu = \sum_i w_{t,i}/K$ be the average weight. We introduce a constant $\ell = \min\{\mu - \delta - 1, \min_i w_{T,i}\}$, which satisfies $w_{T,i} \geq \ell$ for all i at the initial time T. Now assume that $w_{t,i} \geq \ell$ holds for some t. If $x_{t-1,i} = 0$, then $w_{t-1,i} = w_{t,i} + p_{t-1,i} - x_{t-1,i} \geq w_{t,i} \geq \ell$. If $x_{t-1,i} = 1$, then $w_{t-1,i} = w_{t,i} + p_{t-1,i} - x_{t-1,i} \geq \max_j(w_{t,j} + p_{t-1,j}) - \delta - 1 \geq \mu - \delta - 1 \geq \ell$. Thus, in both cases, we obtain $w_{t-1,i} \geq \ell$, and by induction the inequality holds for all t. This establishes a uniform lower bound ℓ for each component of $\boldsymbol{w}_t$. Since a lower bound exists and the total weight $\sum_i w_{t,i}$ is conserved, every component of $\boldsymbol{w}_t$ must remain bounded from above as well, and therefore $\boldsymbol{w}_t$ is bounded. $\square$

Proposition 2. *The discrepancy between the averaged probability vector $\langle \boldsymbol{p}_\tau \rangle_{\tau=0}^{T-1}$ and the empirical distribution $\langle \boldsymbol{x}_\tau \rangle_{\tau=0}^{T-1}$ of samples generated by the herding-based denoising converges to zero as $T \to \infty$ with convergence rate $O(T^{-1})$, where $\langle \cdot \rangle_{\tau=0}^{T-1}$ denotes the average over $\tau = 0, 1, \ldots, T-1$.*

Proof. From Eq. (10), the discrepancy can be rewritten as

$$\langle \boldsymbol{x}_\tau - \boldsymbol{p}_\tau \rangle_{\tau=0}^{T-1} = \langle \boldsymbol{w}_{\tau+1} - \boldsymbol{w}_\tau \rangle_{\tau=0}^{T-1} = \frac{1}{T}(\boldsymbol{w}_T - \boldsymbol{w}_0),$$

which converges to $\boldsymbol{0}$ with rate $O(T^{-1})$ as $T \to \infty$, due to the boundedness of $\boldsymbol{w}_t$. $\square$

Although the asymptotic convergence is guaranteed for any initial weight vector $\boldsymbol{w}_T \in \mathbb{R}^K$, its initialization plays an important role in practical sample generation for finite T.

For uniform diffusion noising, random initialization of $\boldsymbol{w}_T$ from $[0,1]^K$ is justified for the following reason. The stationary distribution of the uniform diffusion Markov chain is the uniform distribution over the sample space $\mathcal{V}$. Therefore, we regard $\boldsymbol{w}_T$ as a weight vector obtained after iterating the herding updates infinitely many times with the uniform probability vector fixed at $\boldsymbol{p}_T = \boldsymbol{1}/K$. After a transient period, each weight $w_{t,i}$ increases by $1/K$ at each iteration but decreases by 1 once every K iterations. As a result, all $w_{T,i}$ remain within an interval of length 1. Accordingly, we choose the uniform distribution on $[0,1]^K$ as a natural initialization.

4. Experiments

We demonstrate the effectiveness of the proposed discrete denoising algorithm through experiments on text and image generation and diffusion-based combinatorial optimization. For clarity of presentation, we reproduce the trained models of UDLM (Schiff et al., 2025) and DIFUSCO (Sun & Yang, 2023), and evaluate the improvements that result solely from introducing the herding algorithm into the reverse denoising process. All experiments are conducted using the official codes provided by the authors. Our algorithm is implemented with only about 20 lines of code modifications, which are included in the Appendix for reproducibility.

Table 2. Evaluation of LM1B text samples by generative PPL (Gen. PPL), entropy, and MAUVE score. The bottom two rows show our results using the UDLM model trained with the code from (Schiff et al., 2025). The herding margin is set to $\delta = 0.0005$ for $T = 1024$ and $\delta = 0.005$ for $T = 128$. Best Gen. PPL values are indicated in bold. [†]Published values from (Schiff et al., 2025).

Method	$T = 1024$			$T = 128$		
	Gen. PPL ($\downarrow$)	Entropy ($\uparrow$)	MAUVE ($\uparrow$)	Gen. PPL ($\downarrow$)	Entropy ($\uparrow$)	MAUVE ($\uparrow$)
AR[†]	–	–	–	67.46	–	–
MDLM[†]	116.80	–	–	120.93	–	–
UDLM[†]	78.22	–	–	79.87	–	–
UDLM (reproduced)	98.75	6.83	0.424	98.33	6.81	0.369
UDLM herding (ours)	**46.71**	6.30	0.648	**60.39**	6.34	0.473

Table 3. Evaluation of CIFAR-10 image samples by FID and IS. The bottom two rows show our results using the UDLM models trained with the code from (Schiff et al., 2025). D-CFG denotes conditional models with $\gamma = 1$. The herding margin is set to $\delta = 0.07$ for $T = 128$ and $\delta = 0.15$ for the other cases. Best values are indicated in bold. [†]Published values from (Austin et al., 2021) and (Schiff et al., 2025).

Method	$T = 1000$		$T = 1024$ (D-CFG)		$T = 128$ (D-CFG)	
	FID ($\downarrow$)	IS ($\uparrow$)	FID ($\downarrow$)	IS ($\uparrow$)	FID ($\downarrow$)	IS ($\uparrow$)
D3PM uniform[†]	51.27	5.99	–	–	–	–
MDLM[†]	33.75	6.74	27.94	7.14	64.09	5.81
UDLM[†]	33.65	6.86	26.70	7.43	30.48	7.30
UDLM (reproduced)	32.64	6.97	24.94	7.41	29.19	7.18
UDLM herding (ours)	**26.39**	**7.40**	**19.20**	**7.79**	**23.95**	**7.51**

4.1. Character-Level Text Generation

We first evaluate the proposed algorithm in character-level text generation, using the text8 dataset (Mahoney, 2011). The vocabulary size is 35 and the output length is fixed to 256 characters. The model architecture follows the UDLM paper (Schiff et al., 2025), which is based on a Transformer with 92.4M parameters. For evaluation, we generate 1,024 samples for each condition and compute the generative perplexity (PPL) based on the GPT-2 Large pretrained model (Radford et al., 2019) and the entropy.

Figure 1 shows the generative PPL and entropy of the generated samples. The herding-based denoising with $\delta = 0$ consistently outperforms stochastic denoising in terms of PPL, achieving up to a $10\times$ improvement for $T = 1024$ inference steps, albeit at the expense of entropy. Increasing the margin δ interpolates between these two extremes, allowing us to improve PPL with a moderate trade-off in entropy.

As DNDM (Chen et al., 2024) provides a deterministic denoising algorithm for masked diffusion, it is informative to contrast its performance on the text8 dataset with ours, although the experimental settings are not exactly identical. DNDM reduces the PPL from 1465.75 of a 1000-step D3PM absorb to 600.02 using 256-step inference. In contrast, our method reduces the PPL from 459.53 of a 1024-step UDLM to 79.16 using 256-step herding, indicating stronger per-formance in both the magnitude of improvement and the achieved generative PPL.

4.2. Word-Level Text Generation

We also evaluate the proposed algorithm in word-level text generation, using the one billion words dataset (LM1B) (Chelba et al., 2013). The vocabulary size is $30,522$ and the output length is 128 tokens. The model architecture follows the UDLM paper (Schiff et al., 2025), which is based on a Transformer with 139M parameters. For evaluation, we generate 1,024 samples for each condition and compute the generative PPL and MAUVE score (Pillutla et al., 2021) based on the GPT-2 Large pretrained model (Radford et al., 2019) as well as the entropy.

Table 2 shows the evaluation results. In our reproduction experiment, the UDLM model trained using wrapped LM1B sequences performs worse than the published results (Schiff et al., 2025). Nevertheless, the herding-based denoising sub-stantially reduces the PPL compared to stochastic denoising for both $T = 1024$ and $T = 128$ steps, outperforming the published UDLM (uniform) and MDLM (masked) results and even the autoregressive (AR) model. We also observe slightly lower entropy values and improved MAUVE scores. Overall, the proposed algorithm significantly improves the coherence and linguistic quality of the generated text while affecting the distributional metrics only minimally. It should

Table 4. Evaluation of TSP solutions obtained by DIFUSCO. The tour lengths and performance gaps, averaged over all instances, are shown for the TSP-500, TSP-1000, and TSP-10000 datasets. The performance gap is computed relative to the best-known baselines[†] (Fu et al., 2021). Greedy and Sampling denote the generation strategies based on one sample with $T = 50$ steps and 16 samples with $T = 10$ steps, respectively. 2-opt indicates post-processing using the 2-opt local search heuristic. The herding margin is set to $\delta = 0.2$. Best values for each strategy are indicated in bold. [‡] Published values of DIFUSCO (Sun & Yang, 2023).

Strategy	Refinement	TSP-500		TSP-1000		TSP-10000	
		Length (↓)	Gap (↓)	Length (↓)	Gap (↓)	Length (↓)	Gap (↓)
Best-known[†]		16.55	0.00%	23.12	0.00%	71.78	0.00%
Greedy[‡]		18.35	–	26.14	–	98.15	–
Greedy (reproduced)		18.24	10.24%	25.61	10.80%	97.49	35.82%
Greedy + herding (ours)		18.06	9.14%	25.94	12.21%	93.18	29.82%
Greedy[‡]	2-opt	16.80	–	23.56	–	73.99	–
Greedy (reproduced)	2-opt	16.81	1.63%	**23.56**	**1.91**%	73.94	3.01%
Greedy + herding (ours)	2-opt	**16.73**	**1.12**%	23.58	2.01%	**73.31**	**2.13**%
Sampling[‡]		17.23	–	25.19	–	95.52	–
Sampling (reproduced)		17.16	3.69%	24.69	6.80%	94.65	31.86%
Sampling + herding (ours)		17.07	3.19%	24.77	7.16%	92.29	28.58%
Sampling[‡]	2-opt	16.65	–	23.45	–	73.89	–
Sampling (reproduced)	2-opt	16.65	0.66%	23.38	1.13%	73.99	3.08%
Sampling + herding (ours)	2-opt	**16.63**	**0.52**%	**23.37**	**1.11**%	73.56	**2.49**%

also be noted that, unlike stochastic denoising, the generative PPL improves as the number of inference steps T increases, which suggests that iterative refinement is enhanced through sample generation following the reverse Markov chain more faithfully.

4.3. Categorical-Valued Image Generation

We apply our herding-based algorithm to image generation, following the experiments of D3PM (Austin et al., 2021) and UDLM (Schiff et al., 2025). While pixel values are ordinal by nature, they are treated as discrete categorical variables in this experiment. Specifically, the forward uniform noising process assigns equal probability to transitions to neighboring and distant pixel values alike, which results in suboptimal performance compared to other noising processes. Accordingly, this experiment is intended only to demonstrate the improvements due to the derandomization.

We use the discretized CIFAR-10 dataset (Krizhevsky & Hinton, 2009), where the pixel intensities of the 32×32 RGB images are discretized to 8-bit integer values. The model is based on a U-Net architecture (Ronneberger et al., 2015), which is identical to that used in (Schiff et al., 2025). For evaluation, we computed Fréchet inception distance (FID) (Heusel et al., 2017) and inception score (IS) (Salimans et al., 2016) on 50,000 generated samples, using the PyTorch evaluation code [1] with a pretrained Inception v3 model (Szegedy et al., 2016).

[1]Available at https://github.com/w86763777/pytorch-image-generation-metrics

From the evaluation results in Table 3, we confirm that the trained unconditional UDLM model and conditional UDLM model (D-CFG with $\gamma = 1$) reproduce the published values nearly identically. The herding-based denoising consistently outperforms stochastic denoising in both FID and IS across all cases. Remarkably, the herding-based denoising with $T = 128$ steps achieves performance comparable to that of stochastic denoising with $T = 1024$ steps.

4.4. Diffusion-Based Combinatorial Optimization

As an alternative evaluation approach, we apply the herding-based denoising to combinatorial optimization. DIFUSCO (Sun & Yang, 2023) is a framework for solving combinatorial optimization problems using discrete diffusion models, where discrete diffusion with Bernoulli noise has been shown to perform significantly better than continuous diffusion with Gaussian noise. In the Bernoulli case, the generative process can be derandomized using the herding-based algorithm, and its performance can be quantitatively evaluated using the objective values of the generated solutions.

We conduct experiments on traveling salesman problems (TSPs), following the experimental settings in Sun & Yang (2023). Solutions are obtained from the discrete diffusion model with binary variables ($K = 2$), where the prediction model is based on anisotropic graph neural networks (AGNNs). We use the official code and pretrained checkpoints provided by the authors (Sun & Yang, 2023) to evaluate the improvement by the herding-based denoising.

Table 4 shows the average objective values of the generated

TSP solutions. These results show that our experiments reproduce the published values very closely. The proposed herding-based denoising outperforms stochastic denoising in almost all cases.

5. Discussion

The proposed denoising algorithm defines a mapping from the initial states $(\boldsymbol{x}_T^{(1:L)}, \boldsymbol{w}_T^{(1:L)})$ to their final states $(\boldsymbol{x}_0^{(1:L)}, \boldsymbol{w}_0^{(1:L)})$, which induces a flow of probability mass within the discrete state space $\mathcal{V}^L$. Each initial state encodes a generated sample $\boldsymbol{x}_0^{(1:L)}$, whose required precision is determined by the granularity of the piecewise isometric mapping. As our approach derandomizes Markov chains, the probability flow can be guided by the transition probabilities, as demonstrated in the numerical experiments with the conditional UDLM model (D-CFG) on the CIFAR-10 dataset. Our algorithm could also be combined with existing inference-time scaling methods (Kit et al., 2025; Singhal et al., 2025; Pani et al., 2025) by updating the particles deterministically.

The initial weight vectors $\boldsymbol{w}_T$ are currently drawn randomly from the uniform distribution on $[0,1]^K$, which is a natural starting point as explained in Section 3.2. Scaling the initial weights can influence the balance between the determinism and randomness of the algorithm. Moreover, there is no reason that all time steps should be treated equally in the accumulation of the discrepancy between the empirical and the target distributions. It is therefore natural to consider weighted accumulation, which has close connections to the Frank–Wolfe algorithm (Bach et al., 2012) and Bayesian quadrature (Huszár & Duvenaud, 2012). The weight scheduling can be an important factor in denoising performance.

We found that introducing delayed switching in the herding algorithm is effective in improving the quality of generated samples. However, it is not clear how to set the appropriate switching margin δ, which can depend on the model, the dataset, the vocabulary size K, and the number of inference steps T. Alternative mechanisms for delayed switching are also possible. Further theoretical and numerical studies will be necessary to clarify the role of delayed switching in deterministic denoising.

These observations on weighted accumulation and delayed switching naturally lead to the idea of continuous-time formulation. Discrete diffusion models formulated as continuous-time Markov chains (CTMCs) have been studied extensively (Campbell et al., 2022; Lou et al., 2024), and can also be naturally derandomized using the continuous-time variant of the herding system, which exhibits chaotic billiard dynamics (Suzuki, 2015). As an alternative, chaotic billiard sampling (Suzuki et al., 2013; Lee et al., 2025),

which can be extended to multi-state spin models such as the Potts model (Suzuki, 2013), can also be applied. Compared with the CTMC formulation, these derandomized billiard systems require computing the exact state transition times induced by collisions with the billiard boundary, which is more time-consuming. Since the exact computation may not contribute to the performance, a time-discretized version, closely related to our proposed algorithm, is considered more practical. Nonetheless, continuous-time formulations may offer advantages in dedicated hardware implementations, as demonstrated in the integrated-circuit implementation of chaotic Boltzmann machines (Yamaguchi et al., 2019). This research direction is related to acceleration methods for the reverse processes of CTMCs (Park et al., 2025; Ren et al., 2025b).

As mentioned in the introduction, we focused on the uniform diffusion process in this paper. However, the proposed herding-based denoising algorithm can also be applied to other Markov-chain formulations of discrete diffusion models. While in (Chen et al., 2024) the state transition times from the absorbing state are externally specified, the proposed deterministic algorithm may allow the transition times to be determined by the system itself according to the transition probability, which is conceptually related to an adaptive sampler for masked diffusion (Ben-Hamu et al., 2025). As the remasking mechanism, as employed in LLaDA (Nie et al., 2025; Zhu et al., 2025), makes the reverse process of masked diffusion more dynamic and enables iterative refinement, our approach could be applied to derandomize the remasking process to further enhance the refinement.

The capability of the herding algorithm is not limited to sampling from categorical distributions. By appropriately designing feature functions, additional constraints, such as symmetry, can be incorporated into the generated samples, which may further enhance the applicability of discrete diffusion models. Although we presented a fully deterministic algorithm, it is also possible to combine the method with probabilistic sampling. Since the boundedness of the weight vector is essential, monitoring its dynamics may help improve the performance of stochastic denoising, as in bounded-error sampling (Yamashita & Suzuki, 2019).

In summary, we presented a novel strategy for derandomizing the reverse denoising process of discrete diffusion models, based on the herding algorithm with weakly chaotic dynamics. Our results provide a positive answer to the key question of whether deterministic denoising can also be effective in discrete diffusion models. As our proposed algorithm is minimal, there are many directions for further improvement. We expect that our approach will contribute to enhancing the significance of discrete diffusion in generative modeling.

Impact Statement

This paper proposes a novel derandomization approach for discrete diffusion models. While generative models may involve broad societal impacts depending on their use, we do not foresee any specific negative societal impacts arising directly from our contribution.

References

Adler, R. L., Kitchens, B. P., Martens, M., Tresser, C. P., and Wu, C. W. The mathematics of halftoning. *IBM Journal of Research and Development*, 47(1):5–15, 2003. doi:10.1147/rd.471.0005.

Aihara, K. and Suzuki, H. Theory of hybrid dynamical systems and its applications to biological and medical systems. *Philosophical Transactions of the Royal Society A*, 368(1930):4893–4914, 2010. doi:10.1098/rsta.2010.0237.

Austin, J., Johnson, D. D., Ho, J., Tarlow, D., and van den Berg, R. Structured denoising diffusion models in discrete state-spaces. In *Proceedings of the Thirty-Fifth Annual Conference on Neural Information Processing Systems*, number 1376, pp. 17981–17993, 2021.

Bach, F., Lacoste-Julien, S., and Obozinski, G. On the equivalence between herding and conditional gradient algorithms. In *Proceedings of the Twenty-Ninth International Conference on Machine Learning*, 2012.

Ben-Hamu, H., Gat, I., Severo, D., Nolte, N., and Karrer, B. Accelerated sampling from masked diffusion models via entropy bounded unmasking. In *Proceedings of the Thirty-Ninth Annual Conference on Neural Information Processing Systems*, 2025.

Bengio, Y., Ducharme, R., Vincent, P., and Janvin, C. A neural probabilistic language model. *Journal of Machine Learning Research*, 3:1137–1155, 2003.

Campbell, A., Benton, J., De Bortoli, V., Rainforth, T., Deligiannidis, G., and Doucet, A. A continuous time framework for discrete denoising models. In *Proceedings of the Thirty-Sixth Annual Conference on Neural Information Processing Systems*, number 2049, pp. 28266–28279, 2022.

Chelba, C., Mikolov, T., Schuster, M., Ge, Q., Brants, T., Koehn, P., and Robinson, T. One billion word benchmark for measuring progress in statistical language modeling, 2013. URL https://arxiv.org/abs/1312.3005.

Chen, T., Zhang, R., and Hinton, G. Analog bits: Generating discrete data using diffusion models with self-conditioning. In *Proceedings of the Eleventh International Conference on Learning Representations*, 2023.

Chen, Z., Yuan, H., Li, Y., Kou, Y., Zhang, J., and Gu, Q. Fast sampling via discrete non-Markov diffusion models with predetermined transition time. In *Proceedings of the Thirty-Eighth Annual Conference on Neural Information Processing Systems*, 2024.

Fu, Z.-H., Qiu, K.-B., and Zha, H. Generalize a small pre-trained model to arbitrarily large TSP instances. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.

Goetz, A. Dynamics of piecewise isometries. *Illinois Journal of Mathematics*, 44(3):465–478, 2000. doi:10.1215/ijm/1256060408.

Goetz, A. Piecewise isometries — an emerging area of dynamical systems. In Grabner, P. and Woess, W. (eds.), *Fractals in Graz 2001*, pp. 135–144, Basel, 2003. Birkhäuser Basel. doi:10.1007/978-3-0348-8014-5_4.

Harvey, N. and Samadi, S. Near-optimal herding. In *Proceedings of The 27th Conference on Learning Theory*, volume 35 of *Proceedings of Machine Learning Research*, pp. 1165–1182, 2014.

Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Proceedings of the Thirty-First Annual Conference on Neural Information Processing Systems*, pp. 6629–6640, 2017.

Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. In *Proceedings of the Thirty-Fourth Annual Conference on Neural Information Processing Systems*, pp. 6840–6851, 2020.

Hoogeboom, E., Nielsen, D., Jaini, P., Forré, P., and Welling, M. Argmax flows and multinomial diffusion: learning categorical distributions. In *Proceedings of the Thirty-Fifth Annual Conference on Neural Information Processing Systems*, number 953, 2021.

Huszár, F. and Duvenaud, D. Optimally-weighted herding is bayesian quadrature. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, 2012.

Inose, H. and Yasuda, Y. A unity bit coding method by negative feedback. *Proceedings of the IEEE*, 51(11):1524–1535, 1963. doi:10.1109/PROC.1963.2622.

Inose, H., Yasuda, Y., and Murakami, J. A telemetering system by code modulation — Δ-Σ modulation. *IRE Transactions on Space Electronics and Telemetry*, SET-8(3):204–209, 1962. doi:10.1109/IRET-SET.1962.5008839.

Kit, L. C., Jeha, P., Frellsen, J., Lio, P., Albergo, M. S., and Vargas, F. Debiasing guidance for discrete diffusion

with sequential Monte Carlo. In *Proceedings of the Workshop Frontiers in Probabilistic Inference: Learning meets Sampling*, 2025.

Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. URL https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf.

Lee, K., Chowdhury, S., and Camsari, K. Y. Noise-augmented chaotic Ising machines for combinatorial optimization and sampling. *Communications Physics*, 8(35), 2025. doi:10.1038/s42005-025-01945-1.

Liang, Y., Huang, R., Lai, L., Shroff, N., and Liang, Y. Absorb and converge: Provable convergence guarantee for absorbing discrete diffusion models. In *Proceedings of the Thirty-Ninth Annual Conference on Neural Information Processing Systems*, 2025a.

Liang, Y., Liang, Y., Lai, L., and Shroff, N. Discrete diffusion models: Novel analysis and new sampler guarantees. In *Proceedings of the Thirty-Ninth Annual Conference on Neural Information Processing Systems*, 2025b.

Lou, A., Meng, C., and Ermon, S. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Proceedings of the Forty-First International Conference on Machine Learning*, number 1333, 2024.

Mahoney, M. text8 dataset. http://mattmahoney.net/dc/text8.zip, 2011.

Nie, S., Zhu, F., You, Z., Zhang, X., Ou, J., Hu, J., Zhou, J., Lin, Y., Wen, J.-R., and Li, C. Large language diffusion models. In *The Thirty-Ninth Annual Conference on Neural Information Processing Systems*, 2025.

Pani, C., Ou, Z., and Li, Y. Test-time alignment of discrete diffusion models with sequential Monte Carlo. In *Proceedings of the Second Workshop on Test-Time Adaptation: Putting Updates to the Test!*, 2025.

Park, Y.-H., Lai, C.-H., Hayakawa, S., Takida, Y., and Mitsufuji, Y. Jump your steps: Optimizing sampling schedule of discrete diffusion models. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025.

Pillutla, K., Swayamdipta, S., Zellers, R., Thickstun, J., Welleck, S., Choi, Y., and Harchaoui, Z. MAUVE: Measuring the gap between neural text and human text using divergence frontiers. In *Proceedings of the Thirty-Fifth Annual Conference on Neural Information Processing Systems*, 2021.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. Technical report, OpenAI, 2019.

Ren, Y., Chen, H., Rotskoff, G. M., and Ying, L. How discrete and continuous diffusion meet: Comprehensive analysis of discrete diffusion models via a stochastic integral framework. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025a.

Ren, Y., Chen, H., Zhu, Y., Guo, W., Chen, Y., Rotskoff, G. M., Tao, M., and Ying, L. Fast solvers for discrete diffusion models: Theory and applications of high-order algorithms. In *Proceedings of the Thirty-Ninth Annual Conference on Neural Information Processing Systems*, 2025b.

Ronneberger, O., Fischer, P., and Brox, T. U-Net: Convolutional networks for biomedical image segmentation. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pp. 234–241, 2015. doi:10.1007/978-3-319-24574-4_28.

Sahoo, S. S., Arriola, M., Gokaslan, A., Marroquin, E. M., Rush, A. M., Schiff, Y., Chiu, J. T., and Kuleshov, V. Simple and effective masked diffusion language models. In *Proceedings of the Thirty-Eighth Annual Conference on Neural Information Processing Systems*, 2024.

Sahoo, S. S., Deschenaux, J., Gokaslan, A., Wang, G., Chiu, J. T., and Kuleshov, V. The diffusion duality. In *Proceedings of the Forty-Second International Conference on Machine Learning*, 2025.

Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training GANs. In *Proceedings of the Thirtieth Annual Conference on Neural Information Processing Systems*, pp. 2234–2242, 2016.

Schiff, Y., Sahoo, S. S., Phung, H., Wang, G., Boshar, S., Dalla-torre, H., de Almeida, B. P., Rush, A. M., Pierrot, T., and Kuleshov, V. Simple guidance mechanisms for discrete diffusion models. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025. Code available at https://github.com/kuleshov-group/discrete-diffusion-guidance.

Shi, J., Han, K., Wang, Z., Doucet, A., and Titsias, M. K. Simplified and generalized masked diffusion for discrete data. In *Proceedings of the Thirty-Eighth Annual Conference on Neural Information Processing Systems*, 2024.

Singhal, R., Horvitz, Z., Teehan, R., Ren, M., Yu, Z., McKeown, K., and Ranganath, R. A general framework for inference-time scaling and steering of diffusion models. In *Proceedings of the Forty-Second International Conference on Machine Learning*, 2025.

Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the Thirty-Second International Conference on Machine Learning*, 2015.

Song, J., Meng, C., and Ermon, S. Denoising diffusion implicit models. In *Proceedings of the Ninth International Conference on Learning Representations*, 2021.

Sun, H., Yu, L., Dai, B., Schuurmans, D., and Dai, H. Score-based continuous-time discrete diffusion models. In *Proceedings of the Eleventh International Conference on Learning Representations*, 2023.

Sun, Z. and Yang, Y. DIFUSCO: Graph-based diffusion solvers for combinatorial optimization. In *Proceedings of the Thirty-Seventh Annual Conference on Neural Information Processing Systems*, 2023. Code available at https://github.com/Edward-Sun/DIFUSCO.

Suzuki, H. Monte Carlo simulation of classical spin models with chaotic billiards. *Physical Review E*, 88:052144, 2013. doi:10.1103/PhysRevE.88.052144.

Suzuki, H. Dynamics of load balancing with constraints. *The European Physical Journal Special Topics*, 223(12): 2631–2635, 2014. doi:10.1140/epjst/e2014-02278-7.

Suzuki, H. Chaotic billiard dynamics for herding. *Nonlinear Theory and Its Applications, IEICE*, 6(4):466–474, 2015. doi:10.1587/nolta.6.466.

Suzuki, H., Imura, J., Horio, Y., and Aihara, K. Chaotic Boltzmann machines. *Scientific Reports*, 3:1610, 2013. doi:10.1038/srep01610.

Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2818–2826, 2016. doi:10.1109/CVPR.2016.308.

Villalobos, P., Ho, A., Sevilla, J., Besiroglu, T., Heim, L., and Hobbhahn, M. Position: Will we run out of data? limits of LLM scaling based on human-generated data. In *Proceedings of the Forty-First International Conference on Machine Learning*, 2024.

von Rütte, D., Fluri, J., Ding, Y., Orvieto, A., Schölkopf, B., and Hofmann, T. Generalized interpolating discrete diffusion. In *Proceedings of the Forty-Second International Conference on Machine Learning*, 2025.

von Rütte, D., Fluri, J., Pooladzandi, O., Schölkopf, B., Hofmann, T., and Orvieto, A. Scaling behavior of discrete diffusion language models, 2025. URL https://arxiv.org/abs/2512.10858.

Welling, M. Herding dynamical weights to learn. In *Proceedings of the Twenty-Sixth International Conference on Machine Learning*, pp. 1121–1128, 2009. doi:10.1145/1553374.1553517.

Welling, M. and Chen, Y. Statistical inference using weak chaos and infinite memory. *Journal of Physics: Conference Series*, 233(1):012005, 2010. doi:10.1088/1742-6596/233/1/012005.

Yamaguchi, M., Katori, Y., Kamimura, D., Tamukoh, H., and Morie, T. A chaotic Boltzmann machine working as a reservoir and its analog VLSI implementation. In *Proceedings of the 2019 International Joint Conference on Neural Networks*, pp. 1–7, 2019. doi:10.1109/IJCNN.2019.8852325.

Yamashita, H. and Suzuki, H. Convergence analysis of herded-Gibbs-type sampling algorithms: effects of weight sharing. *Statistics and Computing*, 29:1035–1053, 2019. doi:10.1007/s11222-019-09852-6.

Zhang, Z., Chen, Z., and Gu, Q. Convergence of score-based discrete diffusion models: A discrete-time analysis. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025.

Zhu, F., Wang, R., Nie, S., Zhang, X., Wu, C., Hu, J., Zhou, J., Chen, J., Lin, Y., Wen, J.-R., and Li, C. LLaDA 1.5: Variance-reduced preference optimization for large language diffusion models, 2025. URL https://arxiv.org/abs/2505.19223.

A. Experimental Details

A.1. Code Modifications for UDLM

We present below the code diff that integrates the proposed herding-based denoising algorithm into the official UDLM codebase (Schiff et al., 2025), for reproducibility of our experimental results in Sections 4.1 to 4.3.

```
--- a/diffusion.py
+++ b/diffusion.py
@@ -142,6 +142,8 @@ class Diffusion(L.LightningModule):

     self.lr = self.config.optim.lr
     self.sampling_eps = config.training.sampling_eps
+    self.deterministic_sampling = getattr(self.config.sampling, 'name', 'gumbel') == 'deterministic'
+    self.deterministic_switch_threshold = getattr(self.config.sampling, 'deterministic_switch_threshold', 1.0)

     self.softplus = torch.nn.Softplus()
     self.neg_infinity = -1_000_000.0
@@ -170,6 +172,22 @@ class Diffusion(L.LightningModule):

     self._validate_configuration()

+  def _sample_categorical_stateful(self, categorical_probs, xt=None):
+    if not self.deterministic_sampling:
+      return _sample_categorical(categorical_probs)
+    self.errors += categorical_probs.to(self.errors.dtype)
+    xs = self._delayed_argmax(xt)
+    self.errors -= F.one_hot(xs, num_classes=self.vocab_size).to(self.errors.dtype)
+    return xs
+
+  def _delayed_argmax(self, xt):
+    if xt is None:
+      return self.errors.argmax(dim=-1)
+    errors_xmax, xmax = self.errors.max(dim=-1)
+    errors_xt = torch.gather(self.errors, -1, xt.unsqueeze(-1)).squeeze(-1)
+    delay_flag = (errors_xmax - errors_xt) < self.deterministic_switch_threshold
+    return torch.where(delay_flag, xt, xmax)
+
   def _validate_configuration(self):
     assert not (self.change_of_variables
                 and self.importance_sampling)
@@ -1123,6 +1141,9 @@ class Diffusion(L.LightningModule):
       self.config.model.length
     ).to(self.device)

+    if self.deterministic_sampling:
+      self.errors = torch.rand(*xt.shape, self.vocab_size, device=xt.device, dtype=torch.float32)
+
     timesteps = torch.linspace(
       1, eps, self.config.sampling.steps + 1, device=self.device)
     dt = (1 - eps) / self.config.sampling.steps
@@ -1243,7 +1264,7 @@ class Diffusion(L.LightningModule):
         f"Diffusion type {self.diffusion} not implemented.")

     # Sample from posterior
-    xs = _sample_categorical(q_xs)
+    xs = self._sample_categorical_stateful(q_xs, xt)
     if self.diffusion == 'absorbing_state':
       copy_flag = (xt != self.mask_index).to(torch.bool)
       q_xs[copy_flag] = 0.0
@@ -1331,7 +1352,7 @@ class Diffusion(L.LightningModule):
         f"Diffusion type {self.diffusion} not implemented.")

     # Sample from posterior
-    xs = _sample_categorical(q_xs)
+    xs = self._sample_categorical_stateful(q_xs, xt)
     if self.diffusion == 'absorbing_state':
       copy_flag = (xt != self.mask_index).to(torch.bool)
       q_xs[copy_flag] = 0.0
@@ -1449,7 +1470,7 @@ class Diffusion(L.LightningModule):

     guided_probs = guided_log_probs.softmax(dim=-1)
     # Sample from guided posterior
-    xs = _sample_categorical(guided_probs)
+    xs = self._sample_categorical_stateful(guided_probs, xt)
     if self.diffusion == 'absorbing_state':
       xs = torch.where(copy_flag.to(bool), xt, xs)
     return xs, guided_probs, {'log_x_theta': log_x_theta,
@@ -1561,7 +1582,7 @@ class Diffusion(L.LightningModule):
       raise NotImplementedError(
         f"Diffusion type {self.diffusion} not implemented.")

-    xs = _sample_categorical(guided_probs)
+    xs = self._sample_categorical_stateful(guided_probs, xt)
     if self.diffusion == 'absorbing_state':
       xs = torch.where(copy_flag, xt, xs)
```

A.2. Code Modifications for DIFUSCO

We present below the code diff that replaces the original denoising algorithm with the proposed herding-based algorithm in the official DIFUSCO codebase (Sun & Yang, 2023), for reproducibility of our experimental results in Section 4.4.

```
--- a/pl_tsp_model.py
+++ b/pl_tsp_model.py
@@ -134,8 +134,19 @@
         else:
            x0_pred_prob = x0_pred.reshape((1, points.shape[0], -1, 2)).softmax(dim=-1)

-        xt = self.categorical_posterior(target_t, t, x0_pred_prob, xt)
-        return xt
+        prob = self.categorical_posterior(target_t, t, x0_pred_prob, xt)
+        self.wt += prob.to(self.wt.dtype)
+        xs = self._delayed_argmax(xt)
+        self.wt -= F.one_hot(xs, num_classes=2).to(self.wt.dtype)
+        return xs
+
+    def _delayed_argmax(self, xt):
+        if xt is None:
+            return self.wt.argmax(dim=-1)
+        wp_max, x_max = self.wt.max(dim=-1)
+        wp_stay = torch.gather(self.wt, -1, xt.unsqueeze(-1)).squeeze(-1)
+        delay_flag = (wp_max - wp_stay) < self.args.ddd_margin # margin from run-time option
+        return torch.where(delay_flag, xt, x_max)

    def gaussian_denoise_step(self, points, xt, t, device, edge_index=None, target_t=None):
        with torch.no_grad():
@@ -202,6 +213,8 @@
        steps = self.args.inference_diffusion_steps
        time_schedule = InferenceSchedule(inference_schedule=self.args.inference_schedule,
                                          T=self.diffusion.T, inference_T=steps)
+
+        self.wt = torch.rand(*xt.shape, 2).to(device) # initialize

        # Diffusion iterations
        for i in range(steps):

--- a/pl_meta_model.py
+++ b/pl_meta_model.py
@@ -136,14 +136,8 @@

        sum_x_t_target_prob += x_t_source_prob_new[..., 1] * x0_pred_prob[..., 1]

-        if target_t > 0:
-            xt = torch.bernoulli(sum_x_t_target_prob.clamp(0, 1))
-        else:
-            xt = sum_x_t_target_prob.clamp(min=0)
-
-        if self.sparse:
-            xt = xt.reshape(-1)
-        return xt
+        prob = sum_x_t_target_prob.clamp(0, 1).reshape(-1)
+        return torch.stack([1-prob, prob], dim=-1)

    def gaussian_posterior(self, target_t, t, pred, xt):
        """Sample (or deterministically denoise) from the Gaussian posterior for a given time step.
```

B. Random Text Samples

Below we present the first eight outputs for each of the stochastic and deterministic denoising algorithms, generated by the same UDLM model trained on the LM1B dataset with $T = 1024$ steps (Table 2).

B.1. Stochastic Denoising

[CLS] lukewarm to crash the world's $ 3 billion (£1. 4 billion) party. [CLS] it has helped to quench anger of the regular customers who long ago bought exclusive movies on the web or through itunes. [CLS] nokia has several partnerships with manufacturers to prevent the mobile phone markets from dying and plans to market several models to only the end sellers by 2010. [CLS] " i am addicted to pain relief. [CLS] do - - vote for app print / audio. [CLS] john s. senior of the university of washington center, seattle, predicted that the previous projection of 2. 01 million jobs this year were created in 2006. [CLS]

[CLS] suffering or the public interest, " the minister told reporters. [CLS] in a rare move in mid - eastern virginia, stevens edged him out of warm washington politics. [CLS] between celebrating with a journey to boot the pumas have become victims of a new year of joy. [CLS] san diego fell to 6 - 0 in 25 of the 22, 000 seats at azteca stadium. [CLS] by the time they had arrived at 3. 9am a riot ensued to which no one would respond. [CLS] scandinavian investors cannot replace workers after their boards lose more than 6, 500 heads every week to other large single ventures, light of day, and today [CLS]

[CLS]iawski, i koch (a barquierja), steven, (espanola aurth), gilbert, juan jean - baptiste (fyzi kosian), dujellschhofer, salim he - nazarzi, korbowu, gousset i yves heymans, pierre glover, isaiah, givat, beauxmele, plygielewski, racec, gedizzola,

camoranesi, m bijzyrafors, b czemolo. [CLS] even so, rain began in the morning, left the field and put a [CLS]

[CLS] " it was inevitably here. [CLS] dr. h. richard hurley of the university of texas injected alexander with a tiny drug histoxipturably k to repair stuttering muscles ahead of his procedure. [CLS] face the press, so why's more? [CLS] as an outsider, oussi is still led by his armed yet steady partner. [CLS] cano drilled a solo homer in the 11th, striking the wrist of rodriguez to juice abdullah's duel with adam jones, the other in season 3 debut. [CLS] meteorologist rossio said as they tested the lake floor, temperatures found themselves hot, reaching sixc. [CLS] [CLS]

[CLS]. [CLS] geithner also exaggerates her comments on leading the united states out of a financial crisis that has had a huge negative impact on the nation's economy, and is on the currency's creation as well as on other nations that have given power to their governments. [CLS] but their response demonstrated weakly that others might preach about direness, and some looked appropriate. [CLS] boiona de gomez of the left (left) did the work until her son lourdes, and son hilda of salvador, both grew up. [CLS] avery has had his tools running through 16 - inning bats and will have 385 at - bats [CLS]

[CLS] the department of justice ; any of the actions needed may be called for the suspension of operations of any portion of the above guarana series due (notes terminate or revoke the note) or be required to transport the vehicle in need ; public web site resources available at website. usw. gov / aftol / increasest and ldcpp for specific. [CLS] how will his dog mane affect people? [CLS] nasa engineers said the name has been up since the 2004, considering an expected number of launches last year from the 2001. [CLS] it is also the way the latest compilation is a mistake. [CLS] christakis, [CLS]

[CLS] brownback partners are the final participants in a half - day meeting. [CLS] unlike avid house - hunters who are more inclined to take videos of quality gadgets like web mani on toys. com, mr o'brien says regular updates on controversial tools or full encompasses of predators and a multitude of natural disasters dog other online markets. [CLS] dawn brancheau had explicitly cycled into the river seine with the remains of her dead husband. [CLS] i wonder why we do not run older companies rootcare problems. [CLS] the report showed americans are still split over u. s. economy's slow pace, deflation levels and unemployment. [CLS] a [CLS]

[CLS] the discussion. - consolidatedtable conclusions both as a historical and comprehensive transaction. sales, revenue and revenues for moving forward in one empty but titled year. been in filings on form s - 8b filed by the sec with the securities and exchange. [CLS] they are normally spent during trials at the court itself. [CLS] djokovic quickly began his resurgence, winning an emphatic ninth game as federer hauled down three break points. [CLS] organised crime, the always, streaked this true. [CLS] " all cities will benefit, " said niamh myung norungon of the university of california, san diego and colleagues at the [CLS]

B.2. Deterministic Denoising

[CLS] renault - nissan alliance, which has been with the white house. [CLS] so check the scrape to tell you what you can... you can. [CLS] the first one you saw, that was called " the doctrine of the doctrine of the faith, what right? " [CLS] who is in damage control? [CLS] " no one wants to shoot at a building next to the supreme court, " he said. [CLS] this has been the only long - running issue - that we turned up for the t20 world championship instead of ashes series and i'm sure that it's going to be the number two decision. [CLS] as the [CLS]

[CLS] get there, i won't pass it up, but i would prefer to be in the country in 2009, " he said. [CLS] so what does this new policy do? [CLS] not even bad, but obama, who had just phoned bush to say it was real, had no need lose himself in a politician's badge of honor. [CLS] the young man's years of abuse started at an age in 1979 when he started praying at the same mosque in birmingham. [CLS] a force spokeswoman said : " while levels of crime at the premises have fallen, it is our responsibility to make sure we can continue these [CLS]

[CLS], she said she and her husband will spend time together after the game in cleveland. [CLS] " we are in progress with confidence, with the strength of our 1752 portfolio and the ability to meet our obligations to the market, " he said. [CLS] all five of the banking groups have promised support for a scheme worth more than £50 billion, 69 % of which was borrowed before the tax change in april. [CLS] his father told the priest that the boy would come out of retirement but would only return home at the end of his life. [CLS] he was also involved in the crash. [CLS] if you happen to wait for the day, [CLS]

[CLS] is any way to stop, speak to me. [CLS] it said that it was channeling shareholders through a " period of difficult trading " that had been far short of expectations. [CLS] " access becomes an issue for the children every time we get out of the school, " the spokesman said. [CLS] " he really is never going to settle down. [CLS] she was the proud parent of three of her children who said it was a nightmare. [CLS] davis said bush's words helped ease her train ride. [CLS] samples of the water, thought to have been drained due to sewage, would be available in the area over three weeks. [CLS] [CLS]

[CLS] near - constant tension in darfur, one of the world's poorest, about 60 miles east of khartoum. [CLS] taken from the gses'perspective, " the urban unemployment rate has dropped 30 percent in the past year, not by more than 7 percent, but only by 5 percent, " the group wrote in the original report, which favored health, safety and financial standards. [CLS] " i'd like you to, " he said. [CLS] " we have the power, one agency at a time... respond to a new law, " he assured congress. [CLS] o. k. the axle [CLS]

[CLS] battle of her own in the us election for the nomination. [CLS] also, officials say 911 callers confirm that the caller is related to a suspect in the vicinity. [CLS] there is no immediate prospect that fiat motors, a large u. s. - owned subsidiary that

770 *competes with the italian automaker fiat, will take over operations of its michigan plant, gm said. [CLS] former lover michelle*
771 *roberts, 31, called the matter a long - distance affair. [CLS] and since foreign troops came to their assistance early in 2007,*
772 *the afghan authorities have faced a growing flood of taliban military forces. [CLS] it offers a lot of features, " said [CLS]*

773 *[CLS] [CLS] and she says : i don't think she's very lucky but today she's feeling better. [CLS] prices are still falling. [CLS]*
774 *this would be microsoft's biggest buy - out exercise ever in yahoo's long history, but executives say microsoft could have time*
775 *to turn it around in just over a year. [CLS] just when the war was under way, u. s. forces were in iraq for the entire year and*
776 *spent nearly four times the u. s. budget for operations in their adopted country in 2003, 2004 and october 2007. [CLS] it also*
777 *includes the vast majority of the big east's major [CLS]*

778 *[CLS] today's economy, new schools and more flexibility, " duncan said. [CLS] and, yes, we were confused. [CLS] the*
779 *next chapter represents an important move toward the right. [CLS] and as it stands today, i hesitate to say this because the*
780 *administration still has a specific intent to put the obama administration's unemployment rate at 7 million to 9. 7 million.*
781 *[CLS] comment : some times he was able to take an on - line gender issue for political gain... maybe he was just asking why.*
782 *[CLS] you know what, we must change if we all want to understand this. [CLS] " there is no time for [CLS]*