
ART for Diffusion Sampling: A Reinforcement Learning Approach to Timestep Scheduling

Yilie Huang*

Wenpin Tang[†]

Xun Yu Zhou[‡]

Abstract

We consider time discretization for score-based diffusion models to generate samples from a learned reverse-time dynamic on a finite grid. Uniform and hand-crafted grids can be suboptimal given a budget on the number of time steps. We introduce Adaptive Reparameterized Time (ART), which controls the clock speed of a reparameterized time variable to redistribute computation along the sampling trajectory while preserving the terminal time, with the objective of minimizing the aggregate Euler discretization error. We derive a randomized companion ART-RL that recasts ART as a continuous-time reinforcement learning problem with Gaussian policies, and prove a two-directional bridge between the two: the deterministic ART optimum lifts to an optimal Gaussian policy, and conversely any optimal Gaussian policy must recover the ART control through its mean. This bridge turns continuous-time actor-critic learning into a principled, rather than heuristic, route to the deterministic timestep optimum. Within the official EDM pipeline, ART-RL improves FID on CIFAR-10 across a wide range of budgets; after one-time offline training, the distilled deterministic schedule transfers without retraining to AFHQv2, FFHQ, and ImageNet at no extra inference cost.

1 Introduction

Diffusion models [6, 31, 32] are a family of generative models that create samples from unknown target distributions, and underpin recent advances in text-to-image creation [26, 27], text-to-video generation [25, 29, 5], and diffusion large language models [19, 24].

A diffusion model consists of two steps: learn the target score function by training on given samples (pretraining), and generate new samples from the learned model (inference). The focus of this paper is on diffusion sampling/inference, where a key numerical choice is the timestep grid used to discretize the learned model. Standard samplers often use uniform or hand-crafted timesteps [30, 32, 18, 2], while recent work also studies optimized fixed schedules [28]. Here our purpose is to provide a control-theoretic framework, which allows strategic and systematic time discretization in diffusion sampling in order to minimize the aggregate error when applying the Euler scheme. The main contribution of this paper is as follows:

- *Formulation and methodology:* We propose an optimal control framework – *Adaptive Reparameterized Time* (ART), to schedule timesteps for diffusion sampling. The idea is to treat the speed of the diffusion sampler as control to reparameterize time and adaptively redistribute computation along the sampling trajectory. To solve the ART problem, we

*Department of Applied Mathematics, The Hong Kong Polytechnic University, Hong Kong. Email: yilie.huang@polyu.edu.hk.

[†]Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027, USA. Email: wt2319@columbia.edu.

[‡]Department of Industrial Engineering and Operations Research and Data Science Institute, Columbia University, New York, NY 10027, USA. Email: xz2574@columbia.edu.

propose a continuous-time reinforcement learning (CTRL) approach – ART-RL, premised upon the recent development in Wang et al. [36], Jia and Zhou [15, 16].

- *Theory*: We establish a two-directional bridge between ART and ART-RL: *sufficiency* shows that the deterministic ART optimum lifts to an optimal Gaussian policy whose mean is exactly the ART control, and *necessity* shows that any optimal Gaussian policy must recover the ART control through its mean. This justifies parameterizing the actor as a Gaussian mean and deploying it deterministically with rigorous guarantees, and yields actor–critic updates for learning adaptive time schedules.
- *Experiments*: Our ART-RL schedule consistently outperforms Uniform, the DPM-Solver log-SNR grid [23], and EDM schedules [18] in a matched CIFAR–10 EDM evaluation. It improves FID across a broad range of sampling time budgets, with large gains at low and moderate budgets and a consistent advantage at the largest budget.
- *Generalization and amortization*: ART-RL separates offline schedule learning from inference-time sampling. After training, the stochastic policy is distilled into a deterministic time-only grid that can be reused as a drop-in schedule across timestep counts and transferred across various datasets without retraining.

To the best of our knowledge, this is the first control/RL formulation of diffusion timestep scheduling through adaptive time reparameterization, and the first to provide a two-directional bridge between a deterministic timestep optimum and an actor–critic learning procedure. The proposed method is data-driven and exhibits robust performance with reusable, amortized transfer.

Relevant literature: Diffusion models as generative tools were first proposed by Ho et al. [6] (DDPM) and Song et al. [30] (DDIM) in the discrete setting. The pioneering work of Song et al. [32] introduces a continuous-time formulation of diffusion models, providing a unified treatment that encompasses earlier discrete-time models. Various sampling methods have been proposed for diffusion inference, including predictor-corrector sampler [32], exponential integrator [38], higher-order solvers [39, 37], and optimized fixed timestep schedules [28]. See also Lee et al. [20], Chen et al. [2, 3], Li et al. [22], Benton et al. [1], Li and Yan [21], Huang et al. [7] for convergence analysis of diffusion models (with either uniform or hand-crafted timesteps).

Continuous-time reinforcement learning (CTRL) was first formulated by Wang et al. [36], where exploration is modeled via stochastic relaxed control to capture the trial-and-error nature of reinforcement learning. Subsequent work developed a model-free theoretical foundation via martingale methods [15–17, 35], established performance guarantees [8, 12, 13], and studied policy optimization [40]. The CTRL framework has also been applied to training and fine-tuning diffusion models in generative AI [4, 10, 41, 42].

Organization: Section 2 reviews diffusion models and introduces the ART control formulation. Section 3 describes how CTRL is used to solve ART, leading to the ART-RL approach. Section 4 presents the ART-RL algorithm, followed by empirical results in Section 5. Section 6 concludes. An extended ART–ART-RL analysis, including all proofs, and additional numerical results appear in the extended version [14].

2 Diffusion Models and ART

2.1 Continuous-Time Score-Based Diffusion Models

We briefly review continuous-time score-based diffusion models; see Tang and Zhao [34] for a recent survey. Given samples from an unknown target distribution, a forward process corrupts data over time $\tau \in [0, T]$ toward a tractable reference distribution, and sampling proceeds by integrating a learned reverse-time process.

Diffusion sampling. A forward process obeys the Itô SDE

$$d\bar{x}(\tau) = -f(\tau)\bar{x}(\tau) d\tau + g(\tau) dw(\tau), \quad (1)$$

where $\tau \in [0, T]$, $\bar{x}(0) \sim p_0$ is the target data distribution, w is a standard Wiener process in $\mathbb{R}^d$, and f and g are coefficient functions. Let p_τ be the law of $\bar{x}(\tau)$ and $S(\tau, x) = \nabla_x \log p_\tau(x)$ be the score.

For sampling, we use the probability flow ODE associated with (1), which shares the same marginals as the reverse-time SDE; see e.g. Tang and Zhao [34, Theorem 5.1]. Denote the backward state by $\tilde{x}(\tau) := \bar{x}(T - \tau)$ with initialization $\tilde{x}(0) \sim p_T$. Replacing S with a trained score model $\hat{S}$ yields

$$\frac{d\tilde{x}(\tau)}{d\tau} = f(T - \tau)\tilde{x}(\tau) + \frac{1}{2}g(T - \tau)^2\hat{S}(T - \tau, \tilde{x}(\tau)), \quad (2)$$

where $\tau \in [0, T]$.

Euler discretization. We integrate (2) on a grid $0 = \tau_0 < \tau_1 < \dots < \tau_K = T$ with step sizes $h_i = \tau_{i+1} - \tau_i$, and write $\tilde{x}_i := \tilde{x}(\tau_i)$. The explicit Euler update is

$$\tilde{x}_{i+1} = \tilde{x}_i + h_i \left[f(T - \tau_i)\tilde{x}_i + \frac{1}{2}g(T - \tau_i)^2\hat{S}(T - \tau_i, \tilde{x}_i) \right]. \quad (3)$$

A uniform grid $\tau_i = iT/K$ is simple but allocates evaluations uniformly even when the reverse dynamics vary along the trajectory. A simple intuition is that early stages, where samples resemble noise, may tolerate coarser steps, while later stages will benefit from finer resolutions. Such considerations motivate adaptive, data-driven schedules that redistribute steps under a given, fixed total time budget T .

2.2 ART: Time Reparameterization via Control

Motivated by the drawback of uniform timesteps in backward sampling discussed above, we introduce a reparameterized sampling clock and formalize adaptive timestep selection in the reverse process as a control problem. The key idea is to replace the uniform progression of original time by an adaptive, controlled progression that can accelerate in some segments and decelerate in others, thereby redistributing computational effort along the trajectory.

To this end, let $\psi : [0, T] \rightarrow \mathbb{R}$ be a continuous time mapping from the new clock t to the original diffusion time τ (i.e., $\tau = \psi(t)$), with $\psi(0) = 0$ and $\psi(T) = T$. On this new clock we write the state as $x(t) := \tilde{x}(\psi(t))$, namely the backward state evaluated at diffusion time $\psi(t)$, with $x(0) \sim p_T$. We define the control as $\theta(t) := \dot{\psi}(t)$, which quantifies the instantaneous rate of change in diffusion time with respect to t and satisfies $\int_0^T \theta(t) dt = \psi(T) - \psi(0) = T$. To keep the formulation general, we do not impose monotonicity of ψ , so θ may take either sign. The monotone time-change case (equivalently, $\theta(t) \geq 0$ almost everywhere) is included as a special case of this formulation. In particular, if we discretize the new clock uniformly as $0 = t_0 < t_1 < \dots < t_K = T$, then the induced original-time grid is $\tau_i := \psi(t_i)$ with nonuniform step sizes $\tau_{i+1} - \tau_i$. Intuitively, $x(\cdot)$ tracks the generative trajectory under the reparameterized time, while $\theta(\cdot)$ determines how quickly the trajectory advances in original time per unit step on the new clock (larger $\theta(t)$ corresponds to faster progression at t). We refer to this time reparameterized sampling framework as *Adaptive Reparameterized Time* (ART).

By the chain rule, ART induces the controlled dynamics

$$\begin{cases} \dot{x}(t) = \theta(t)F(x(t), \psi(t)), & x(0) \sim p_T, \\ \dot{\psi}(t) = \theta(t), & \psi(0) = 0, \psi(T) = T, \end{cases} \quad (4a)$$

$$(4b)$$

where the backward probability-flow vector evaluated at original time $T - \psi$ is

$$F(x, \psi) := f(T - \psi)x + \frac{1}{2}g(T - \psi)^2\hat{S}(T - \psi, x). \quad (5)$$

The constraint $\int_0^T \theta(t) dt = T$ is the time budget allocated by the control.

To assess how Euler behaves on the t -clock, we relate its local approximation error over a step to the curvature of the right-hand side in (4a). We proceed analogously to the Euler discretization in (3). On an interval $t \in [t_i, t_{i+1})$ with stepsize $h_i := t_{i+1} - t_i$, the Euler update uses one control value per step; we denote this step value by θ_i .

Define the one-step Euler error proxy on the t -clock by $E_i := x(t_{i+1}) - \{x(t_i) + h_i\theta_i F(x(t_i), \psi(t_i))\}$.

A Taylor expansion yields $E_i = \frac{h_i^2}{2}\theta_i^2 Q(x(t_i), \psi(t_i)) + O(h_i^3)$, where

$$Q(x, \psi) = A(s, x)B(s, x) - g(s)g'(s)\hat{S}(s, x) - f'(s)x - \frac{1}{2}g(s)^2\partial_s\hat{S}(s, x), \quad (6)$$

with $s := T - \psi$, $A(s, x) := f(s)I_d + \frac{1}{2}g(s)^2\nabla_x\hat{S}(s, x)$, and $B(s, x) := f(s)x + \frac{1}{2}g(s)^2\hat{S}(s, x)$. Thus, to leading order, the magnitude of the Euler local error on step i is proportional to $\theta_i^2|Q(x(t_i), \psi(t_i))|$, where $|\cdot|$ denotes the Euclidean norm throughout. In implementation, $\nabla_x\hat{S}(s, x)$ is not formed; automatic differentiation provides the needed Jacobian–vector product.

This motivates using $|Q(x, \psi)|\theta(t)^2$ as an error-density surrogate for allocating the time-warping rate. With a fixed time budget, we choose $\theta = \theta(\cdot)$ to minimize the overall residual surrogate subject to the constraint $\int_0^T \theta(t) dt = T$, enforced via a Lagrange multiplier $\gamma \in \mathbb{R}$. Equivalently, we maximize the negative Lagrangian cost

$$J^\theta(s, y, \phi) = \mathbb{E} \left[\int_s^T (|Q(x(t), \psi(t))|\theta^2(t) - \gamma\theta(t)) dt + \gamma T \mid x(s) = y, \psi(s) = \phi \right]. \quad (7)$$

Denote the optimal value function to be

$$V(s, y, \phi) = \max_{\theta=\theta(\cdot)} J^\theta(s, y, \phi). \quad (8)$$

In summary, ART recasts timestep allocation as continuous-time control of the time-warping rate θ under (4) and (7). The next section develops an RL-based procedure (ART-RL) to learn θ in a data-driven way.

3 Randomized Control and RL Formulation

The ART problem has no closed-form solution and its HJB is intractable on the high-dimensional state space $x \in \mathbb{R}^d$. We therefore introduce a Gaussian randomization of the control as a technical device that embeds ART into the continuous-time RL framework, enabling actor–critic learning. We call this approach *Adaptive Reparameterized Time via Reinforcement Learning* (ART-RL).

3.1 ART-RL: Auxiliary Problem with Gaussian Policies

We model control randomization with a stochastic policy over the time-warping rate θ . Inspired by entropy-regularized control [43, 11], we consider the Gaussian family indexed by $\lambda \geq 0$:

$$\pi^{(\lambda)}(\cdot \mid t, x, \psi) = \mathcal{N} \left(\mu(t, x, \psi), \frac{\lambda}{|Q(x, \psi)|} \right), \quad (9)$$

where μ is a measurable deterministic policy and Q is defined in (6). Let $\Pi^{(\lambda)}$ be the class of such policies with finite second moments. Since Q is proportional to the Euler local-error proxy, the variance is smaller in stiff regions and larger elsewhere; λ controls the overall noise level without changing the mean. For analysis we assume $|Q(x, \psi)| > 0$ almost surely on compact intervals; in practice we use $|Q| \vee \varepsilon$ for small $\varepsilon > 0$.

Given a policy $\pi^{(\lambda)} \in \Pi^{(\lambda)}$, let $(x^{\pi^{(\lambda)}}(t), \psi^{\pi^{(\lambda)}}(t))_{t \in [0, T]}$ denote the corresponding state process which, according to Wang et al. [36], satisfy

$$\begin{cases} \dot{x}^{\pi^{(\lambda)}}(t) = \int_{\mathbb{R}} \theta F(x^{\pi^{(\lambda)}}(t), \psi^{\pi^{(\lambda)}}(t)) \pi_t^{(\lambda)}(\theta) d\theta, & x^{\pi^{(\lambda)}}(0) = x_0 \sim p_T, \end{cases} \quad (10a)$$

$$\begin{cases} \dot{\psi}^{\pi^{(\lambda)}}(t) = \int_{\mathbb{R}} \theta \pi_t^{(\lambda)}(\theta) d\theta, & \psi^{\pi^{(\lambda)}}(0) = 0, \psi^{\pi^{(\lambda)}}(T) = T. \end{cases} \quad (10b)$$

The corresponding total-time constraint is $\int_0^T \int_{\mathbb{R}} \theta \pi_t^{(\lambda)}(\theta) d\theta dt = T$. The objective function for the auxiliary problem is

$$J^{\pi^{(\lambda)}}(s, y, \phi) = \mathbb{E} \left[\gamma T + \lambda T + \int_s^T \int_{\mathbb{R}} \pi_t^{(\lambda)}(\theta) (|Q(x^{\pi^{(\lambda)}}(t), \psi^{\pi^{(\lambda)}}(t))|\theta^2 - \gamma\theta) d\theta dt \mid x^{\pi^{(\lambda)}}(s) = y, \psi^{\pi^{(\lambda)}}(s) = \phi \right],$$

where the Lagrange multiplier $\gamma \in \mathbb{R}$ enforces the total-time constraint and $\lambda \geq 0$ is the variance parameter from (9). The optimal value function is

$$V^{(\lambda)}(s, y, \phi) = \max_{\pi^{(\lambda)} \in \Pi^{(\lambda)}} J^{\pi^{(\lambda)}}(s, y, \phi). \quad (11)$$

3.2 Relationship between ART and ART-RL

We now establish the relationship between the ART control problem (8) and the ART-RL auxiliary problem (11). By dynamic programming, the optimal value function V for (8) satisfies

$$V_t + \sup_{\theta} \left\{ (V_x^\top F(x, \psi) + V_\psi - \gamma)\theta - |Q(x, \psi)|\theta^2 \right\} = 0, \quad (12)$$

with terminal condition $V(T, x, \psi) = \gamma T$.

Moreover, by Wang et al. [36], the optimal value function $V^{(\lambda)}$ of (11) satisfies the exploratory HJB:

$$V_t^{(\lambda)} + \sup_{\mu} \left\{ (V_x^{(\lambda)\top} F(x, \psi) + V_\psi^{(\lambda)} - \gamma)\mu - |Q(x, \psi)| \left(\mu^2 + \frac{\lambda}{|Q(x, \psi)|} \right) \right\} = 0 \quad (13)$$

with terminal condition $V^{(\lambda)}(T, x, \psi) = (\gamma + \lambda)T$.

The following three results characterize the connection: a value-shift theorem aligning the deterministic and auxiliary HJB equations, followed by sufficiency and necessity theorems that together establish a two-directional bridge. All proofs can be found in the extended version [14].

Theorem 3.1 (Value function shift). *If V is a classical solution of (12), then*

$$V^{(\lambda)}(t, x, \psi) = V(t, x, \psi) + \lambda t$$

is a classical solution of (13).

Theorem 3.2 (Sufficiency: deterministic ART optimum lifts to an optimal Gaussian policy). *Under the conditions of Theorem 3.1, define*

$$\mu^*(t, x, \psi) = \frac{V_x^\top F(x, \psi) + V_\psi - \gamma}{2|Q(x, \psi)|}.$$

Then the Gaussian policy

$$\pi^{(\lambda)*}(\cdot | t, x, \psi) = \mathcal{N}\left(\mu^*(t, x, \psi), \frac{\lambda}{|Q(x, \psi)|}\right)$$

is optimal for the auxiliary problem (11) subject to the dynamics (10). Furthermore, μ^* is an optimal policy for the original problem (8) subject to the dynamics (4).

Theorem 3.3 (Necessity: optimal auxiliary policy recovers the ART control). *Suppose Theorem 3.1 holds, $|Q(x, \psi)| > 0$, and the policy-gap residual $R^{\pi^{(\lambda)}}(t, x, \psi) = |Q(x, \psi)|(\mu^{\pi^{(\lambda)}}(t, x, \psi) - \mu^*(t, x, \psi))^2$ is continuous in (t, x, ψ) . Let $\pi^{(\lambda)*} \in \Pi^{(\lambda)}$ be any Gaussian policy that is optimal for the auxiliary problem (11) at every initial pair (s, y, ϕ) . Then its mean field satisfies*

$$\mu^{\pi^{(\lambda)*}}(t, x, \psi) = \mu^*(t, x, \psi) \quad \text{for every } (t, x, \psi),$$

where μ^* is the optimal ART control of Theorem 3.2. In particular, $\mu^{\pi^{(\lambda)*}}$ is also an optimal control for the original ART problem (8).

Theorems 3.1–3.3 establish a *two-directional* equivalence between ART and ART-RL. Hence the actor in Section 4 can be parameterized as a Gaussian mean and deployed deterministically, with the guarantee that any optimal auxiliary policy recovers the ART control through its mean. This bridge is what makes continuous-time actor–critic learning a principled, rather than heuristic, route to the deterministic timestep problem.

4 ART-RL Actor–Critic Algorithm

Theorems 3.1–3.3 reduce ART to solving the auxiliary ART-RL problem and guarantee that any optimal auxiliary policy mean recovers the deterministic ART optimum.

4.1 Algorithm Design

Our algorithm builds on the continuous-time actor-critic framework of Jia and Zhou [16], adapted to ART-RL with randomized Gaussian control. We use two neural networks NN^{ϑ_c} and NN^{ϑ_a} for the value function and policy mean, and parameterize

$$\hat{V}^{\vartheta_c}(t, x, \psi) = NN^{\vartheta_c}(t, x, \psi) + \lambda t, \quad \hat{\pi}^{\vartheta_a}(\cdot | t, x, \psi) = \mathcal{N}\left(NN^{\vartheta_a}(t, x, \psi), \frac{\lambda}{|Q(x, \psi)|}\right). \quad (14)$$

Let $\theta^{\hat{\pi}^{\vartheta_a}} \sim \hat{\pi}^{\vartheta_a}$ be a Gaussian control process with corresponding (observable) state process $(x^{\theta^{\hat{\pi}^{\vartheta_a}}}(t), \psi^{\theta^{\hat{\pi}^{\vartheta_a}}}(t))$. For brevity, in (15) we write $(x(t), \psi(t), \theta(t)) = (x^{\theta^{\hat{\pi}^{\vartheta_a}}}(t), \psi^{\theta^{\hat{\pi}^{\vartheta_a}}}(t), \theta^{\hat{\pi}^{\vartheta_a}}(t))$. By Jia and Zhou [16], $\hat{V}^{\vartheta_c}$ and $\hat{\pi}^{\vartheta_a}$ satisfy the coupled moment conditions

$$\begin{aligned} \mathbb{E}\left[\int_0^T \frac{\partial NN^{\vartheta_c}(t, x(t), \psi(t))}{\partial \vartheta_c} \left(d\hat{V}^{\vartheta_c}(t, x(t), \psi(t)) - (|Q(x(t), \psi(t))|\theta(t)^2 + \gamma\theta(t))dt\right)\right] &= 0, \\ \mathbb{E}\left[\int_0^T \frac{\partial \log \hat{\pi}^{\vartheta_a}(\theta(t) | t, x, \psi)}{\partial \vartheta_a} \left(d\hat{V}^{\vartheta_c}(t, x(t), \psi(t)) - (|Q(x(t), \psi(t))|\theta(t)^2 + \gamma\theta(t))dt\right)\right] &= 0. \end{aligned} \quad (15)$$

These are standard moment conditions in RL [33, 9]; the parameters ϑ_c and ϑ_a are updated by stochastic approximation.

4.2 Implementable ART-RL Algorithm

At iteration n , a trajectory (x_n, ψ_n, θ_n) is generated under the current policy $\hat{\pi}^{\vartheta_{a,n}}$ on a uniform grid $0 = t_0 < \dots < t_K = T$ with $\Delta t = T/K$, and parameters are updated with learning rate $a_n > 0$. Writing $\hat{V}_k^{\vartheta_{c,n}} := \hat{V}^{\vartheta_{c,n}}(t_k, x_n(t_k), \psi_n(t_k))$ and applying stochastic approximation with Riemann discretization to the moment conditions (15) gives, with $D_{n,k} := \hat{V}_{k+1}^{\vartheta_{c,n}} - \hat{V}_k^{\vartheta_{c,n}} - \gamma_n \theta_n(t_k) \Delta t - |Q(x_n(t_k), \psi_n(t_k))|\theta_n(t_k)^2 \Delta t$, the critic and actor updates

$$\vartheta_{c,n+1} \leftarrow \vartheta_{c,n} + a_n \sum_{k=0}^{K-1} \frac{\partial NN^{\vartheta_{c,n}}}{\partial \vartheta_c}(t_k, x_n(t_k), \psi_n(t_k)) D_{n,k}, \quad (16a)$$

$$\vartheta_{a,n+1} \leftarrow \vartheta_{a,n} + a_n \sum_{k=0}^{K-1} \frac{\partial \log \hat{\pi}_{t_k}^{\vartheta_{a,n}}(\theta_n(t_k))}{\partial \vartheta_a} D_{n,k}, \quad (16b)$$

where $\hat{\pi}_{t_k}^{\vartheta_{a,n}}(\theta) := \hat{\pi}^{\vartheta_{a,n}}(\theta | t_k, x_n(t_k), \psi_n(t_k))$. The Lagrange multiplier is updated by $\gamma_{n+1} \leftarrow \gamma_n + a_n(\psi_n(T) - T)$, where $\psi_n(T) = \psi_n(t_K)$. The full procedure is summarized in Algorithm 1.

Algorithm 1 Time-discretized ART-RL Actor-Critic

```

for  $n = 1$  to  $N$  do
  Set  $k = 0, t = t_k = 0$ , initialize  $(x_n(t_0), \psi_n(t_0))$ 
  while  $t < T$  do
    Sample control  $\theta_n(t_k)$  from the Gaussian policy (14)
    Update the states by the ART dynamics (4)
    Increment time:  $t_{k+1} = t_k + \Delta t, k \leftarrow k + 1$ 
  end while
  Collect trajectory  $\{(t_k, x_n(t_k), \psi_n(t_k), \theta_n(t_k))\}_{k=0}^{K-1}$ 
  Update critic and actor parameters via (16a) and (16b)
  Update the Lagrange multiplier  $\gamma_{n+1}$ 
end for

```

5 Numerical Experiments

We evaluate ART-RL in a one-dimensional known-score setting and in image-generation experiments using the official EDM pipeline [18]. Across CIFAR-10, AFHQv2, FFHQ, and ImageNet, the score




Figure 1: Quantitative overview: ART-RL Pareto-dominates Uniform, DPM-logSNR, and EDM on CIFAR-10 across NFE budgets, and the same distilled grid retains the advantage on AFHQv2, FFHQ, and ImageNet without retraining. DPM-logSNR is the DPM-Solver uniform log-SNR grid.

model, solver, and EDM implementation are kept fixed and only the timesteps are replaced. We compare Uniform, the DPM-Solver log-SNR grid, EDM timesteps, and ART-RL, reporting W_2 on the 1D problem and FID vs. NFE on images; on CIFAR-10 we additionally include AYS [28]. Additional experimental details and results are in the extended version [14].

5.1 One-Dimensional Example with Analytical Score

We first consider a one-dimensional example where the score function is analytical in closed form; hence every performance difference comes from the timestep schedule rather than score-estimation error. The forward process starts from $p_0 = \mathcal{N}(0, 1)$ and follows $dx(t) = \sqrt{2t} dw(t)$ over $T = 3$, so $x(t) \sim \mathcal{N}(0, 1 + t^2)$, $p_T = \mathcal{N}(0, 10)$, and $S(t, x) = -x/(1 + t^2)$. Substituting this score into (5) and (6) gives $F(x, \psi) = -(T - \psi)x/(1 + (T - \psi)^2)$ and $Q(x, \psi) = x/(1 + (T - \psi)^2)^2$.

This controlled setting separates learned from hand-designed schedules: EDM and DPM-logSNR are strong but fixed analytic prescriptions, whereas ART-RL learns the time allocation from the objective. In Figure 1-(a), ART-RL lowers W_2 across all tested K ; for example, at $K = 10$, ART-RL obtains .079 versus .177 (EDM) and .211 (DPM-logSNR), and at $K = 100$, .013 versus .023 and .027.

Distillation diagnostic. Training trajectories concentrate tightly around a smooth mean control, so the stochastic policy can be distilled into a deterministic time-only grid, with two practical benefits. First, distillation removes the per-step cost of actor and Q evaluations, so ART-RL sampling requires no extra computation beyond standard schemes such as Uniform or EDM. Second, it eliminates the residual $\psi(T)$ overshoot/undershoot that occurs when θ is produced by a neural-network actor: normalizing the deterministic increments to sum exactly to T guarantees the induced grid hits the terminal time precisely, improving the numerical fidelity of the discretized probability flow ODE.

5.2 CIFAR-10 under the EDM Pipeline

We first run a solver-aligned Euler ablation to verify the objective-sampler alignment. ART-RL improves FID over Uniform, DPM-logSNR, and EDM across all Euler budgets; for example, at NFE = 30 it attains 3.46 versus 3.95 (DPM-logSNR), 4.21 (EDM), and 85.83 (Uniform).

We report the main image results with the default Heun solver. Heun is the sampler used in the official EDM pipeline and a 2nd-order Runge-Kutta method (one order higher than Euler), which makes it both the standard choice for modern low-NFE deterministic diffusion sampling and the configuration in which EDM reports its strongest FID; running under Heun therefore tests ART-RL against the strongest baseline setup. Although our training surrogate is Euler-based, we assess transfer of the Euler-trained schedule under Heun empirically below. We test $K \in \{2, 5, 7, 10, 18\}$ with Euler only

for the final step update, matching the EDM implementation and giving $\text{NFE} = 2K - 1$. Table 1 shows that ART-RL attains the lowest FID across all baselines. The improvement is most pronounced at small and moderate budgets, and remains present at the largest-budget setting where EDM reported its strongest performance: at $\text{NFE} = 35$, ART-RL obtains 1.82, compared with 1.85 for EDM and 1.89 for DPM-logSNR.

Table 1: CIFAR-10 FID vs. NFE. DPM-logSNR and EDM use fixed analytic schedules, whereas AYS is an optimized fixed reference schedule. At $\text{NFE} = 3$, all three use the same boundary noise levels and therefore yield the same FID; ART-RL instead learns its schedule from data.

NFE	3	9	13	19	35
Uniform	280.29	213.13	191.69	168.87	118.02
AYS	465.83	157.29	86.82	35.30	3.60
DPM-logSNR	465.83	52.29	8.67	2.76	1.89
EDM	465.83	35.54	6.79	2.54	1.85
ART-RL	152.86	32.13	5.44	2.45	1.82

Visual comparisons corroborate the table: at very small budgets ART-RL already produces recognizable objects while EDM and DPM-logSNR remain close to noise.

Statistical significance. At $\text{NFE} = 35$, three matched 50,000-sample runs give ART-RL FIDs 1.82, 1.79, 1.82 versus EDM FIDs 1.85, 1.83, 1.85; ART-RL’s worst run beats EDM’s best. For context, the ART-RL vs. EDM gaps on CIFAR-10 are 313, 3.4, and 1.35 at $\text{NFE} = 3, 9, 13$, respectively; these selected gaps are larger than the seed-to-seed 50k-sample FID standard deviation reported by EDM (≈ 0.02 – 0.05).

5.3 Generalization of the ART-RL Time Schedule

The ART-RL schedule is learned once on CIFAR-10 (1–2 hours on a Colab T4, offline), distilled to a fixed grid, and reused across timestep counts and datasets without retraining. After distillation, inference uses only the precomputed grid, just as Uniform, DPM-logSNR, and EDM do. For the transfer studies below we focus on the strongest baselines (DPM-logSNR and EDM); Uniform and AYS are omitted since both are substantially weaker in the matched CIFAR-10 evaluation (Table 1).

5.3.1 Robustness Across Time Grids via Interpolation and Extrapolation

To test reuse across timestep counts, we take the ART-RL schedule learned on CIFAR-10 at $K = 18$ and construct new grids for $K \in \{4, 9, 12, 15, 20\}$ by log-linear interpolation and extrapolation; EDM and DPM-logSNR use their analytic timestep sequences.

Table 2: CIFAR-10 interpolation/extrapolation FID.					
NFE	7	17	23	29	39
DPM-logSNR	185.63	3.52	2.19	1.94	1.88
EDM	85.80	3.11	2.06	1.88	1.85
ART-RL	33.73	2.57	2.00	1.84	1.82

Table 2 shows ART-RL preserves the Section 5.2 hierarchy across interpolated and extrapolated counts.

5.3.2 Cross-Dataset Transfer: AFHQv2, FFHQ, and ImageNet

We plug the CIFAR-10 ART-RL schedule into the official EDM checkpoints for AFHQv2, FFHQ, and ImageNet, testing a drop-in replacement without dataset-specific retraining.




Figure 2: ImageNet samples across NFEs.

Table 3: Cross-dataset FID using the CIFAR-10 ART-RL schedule.

Dataset	Method	3	9	13	19	35
AFHQv2	DPM-logSNR	375.76	67.64	9.77	3.44	2.15
	EDM	375.76	27.88	7.56	2.99	2.11
	ART-RL	243.48	20.48	6.12	2.85	2.07
FFHQ	DPM-logSNR	466.76	113.87	15.94	5.25	2.66
	EDM	466.76	57.13	15.87	5.26	2.73
	ART-RL	305.97	35.73	11.08	4.31	2.57
ImageNet	DPM-logSNR	437.42	60.48	12.31	4.46	2.66
	EDM	437.42	35.32	8.18	3.68	2.57
	ART-RL	147.21	29.49	7.01	3.62	2.56

ART-RL improves over EDM and DPM-logSNR at low and moderate NFEs across all three datasets (Figure 2 and further visual comparisons in Huang et al. [14]), confirming a reusable rather than dataset-specific parameterization.

6 Conclusions

We introduced ART for timestep allocation in score-based diffusion models. Its main modeling step is to replace direct grid selection by a continuous time change whose rate is a learnable control that redistributes computation under a fixed time budget. ART-RL solves this control problem through an auxiliary randomized formulation: sufficiency and necessity theorems show that the deterministic ART optimum and the mean of any optimal Gaussian auxiliary policy coincide, making actor-critic learning a principled route to a deterministic schedule. Empirically, within matched EDM samplers, the distilled ART-RL grid improves FID over Uniform, DPM-logSNR, and EDM on CIFAR-10, generalizes across step counts, and transfers without retraining to AFHQv2, FFHQ, and ImageNet as a drop-in time grid with no inference overhead.

Future directions include extending ART to stochastic samplers, designing higher-order or solver-aware error surrogates, and identifying settings where state-dependent schedules improve over distilled time-only grids. More broadly, adaptive time reparameterization for diffusion sampling remains nascent, and ART with ART-RL offers a control-theoretic starting point for systematic schedule design in generative diffusion models.

Acknowledgments and Disclosure of Funding

Funding. Yilie Huang acknowledges financial support from the Start-up Fund of The Hong Kong Polytechnic University (Project ID: P0063874). Wenpin Tang is supported by NSF CAREER Award DMS-2538791 and NSF MFAI Award DMS-2602038. Xun Yu Zhou is supported by the Nie Center for Intelligent Asset Management and the Yao Artificial Intelligence Research Fund at Columbia University.

Competing Interests. The authors declare no competing interests.

References

- [1] Joe Benton, Valentin De Bortoli, Arnaud Doucet, and George Deligiannidis. Nearly d -linear convergence bounds for diffusion models via stochastic localization. In *International Conference on Learning Representations*, 2024.
- [2] Hongrui Chen, Holden Lee, and Jianfeng Lu. Improved analysis of score-based generative modeling: User-friendly bounds under minimal smoothness assumptions. In *International Conference on Machine Learning*, pages 4735–4763, 2023.
- [3] Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru R Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *International Conference on Learning Representations*, 2023.
- [4] Xuefeng Gao, Jiale Zha, and Xun Yu Zhou. Reward-directed score-based diffusion models via q-learning. *Journal of Machine Learning Research*, 26(302):1–46, 2025.
- [5] Google. State-of-the-art video and image generation with Veo 2 and Imagen 3. <https://blog.google/technology/google-labs/video-image-generation-update-december-2024/>, 2024. Accessed: 2025-09-17.
- [6] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851, 2020.
- [7] Daniel Zhengyu Huang, Jiaoyang Huang, and Zhengjiang Lin. Convergence analysis of probability flow ODE for score-based generative models. *arXiv preprint arXiv:2404.09730*, 2025.
- [8] Yilie Huang. Continuous-time reinforcement learning for asset–liability management. In *ACM International Conference on AI in Finance*, pages 360–368, 2025.
- [9] Yilie Huang and Xun Yu Zhou. Data-driven exploration for a class of continuous-time indefinite linear–quadratic reinforcement learning problems. *IEEE Transactions on Automatic Control*, 2026. Early Access, doi: 10.1109/TAC.2026.3734857.
- [10] Yilie Huang and Xun Yu Zhou. Amortized guidance for image inpainting with pretrained diffusion models. In *Advances in Neural Information Processing Systems*, 2026. To appear; arXiv:2605.13010.
- [11] Yilie Huang, Yanwei Jia, and Xun Yu Zhou. Achieving mean–variance efficiency by continuous-time reinforcement learning. In *ACM International Conference on AI in Finance*, pages 377–385, 2022.
- [12] Yilie Huang, Yanwei Jia, and Xun Yu Zhou. Sublinear regret for a class of continuous-time linear-quadratic reinforcement learning problems. *SIAM Journal on Control and Optimization*, 63(5):3452–3474, 2025.
- [13] Yilie Huang, Yanwei Jia, and Xun Yu Zhou. Mean–variance portfolio selection by continuous-time reinforcement learning: Algorithms, regret analysis, and empirical study. *Management Science*, 2026. To appear; arXiv:2412.16175.
- [14] Yilie Huang, Wenpin Tang, and Xun Yu Zhou. Adaptive reparametrized time for score-based diffusion sampling, 2026. Extended version, arXiv:2607.02137.
- [15] Yanwei Jia and Xun Yu Zhou. Policy evaluation and temporal-difference learning in continuous time and space: A martingale approach. *Journal of Machine Learning Research*, 23(154):1–55, 2022.
- [16] Yanwei Jia and Xun Yu Zhou. Policy gradient and actor-critic learning in continuous time and space: Theory and algorithms. *Journal of Machine Learning Research*, 23(275):1–50, 2022.

- [17] Yanwei Jia and Xun Yu Zhou. q -learning in continuous time. *Journal of Machine Learning Research*, 24(161):1–61, 2023.
- [18] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems*, volume 35, pages 26565–26577, 2022.
- [19] Samar Khanna, Siddhant Kharbanda, Shufan Li, Harshit Varma, Eric Wang, Sawyer Birnbaum, Ziyang Luo, Yanis Miraoui, Akash Palrecha, and Stefano Ermon. Mercury: Ultra-fast language models based on diffusion. *arXiv preprint arXiv:2506.17298*, 2025.
- [20] Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence for score-based generative modeling with polynomial complexity. In *Advances in Neural Information Processing Systems*, volume 35, pages 22870–22882, 2022.
- [21] Gen Li and Yuling Yan. Adapting to unknown low-dimensional structures in score-based diffusion models. In *Advances in Neural Information Processing Systems*, volume 37, pages 126297–126331, 2024.
- [22] Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. Towards faster non-asymptotic convergence for diffusion-based generative models. In *International Conference on Learning Representations*, 2024.
- [23] 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. In *Advances in Neural Information Processing Systems*, volume 35, pages 5775–5787, 2022.
- [24] Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv preprint arXiv:2502.09992*, 2025.
- [25] OpenAI. Sora: Creating video from text. <https://openai.com/sora>, 2024. Accessed: 2025-09-17.
- [26] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [28] Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align your steps: Optimizing sampling schedules in diffusion models. In *International Conference on Machine Learning*, pages 42947–42975, 2024.
- [29] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyan Hu, Harry Yang, Oron Ashual, and Oran Gafni. Make-A-Video: Text-to-video generation without text-video data. In *International Conference on Learning Representations*, 2023.
- [30] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021.
- [31] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, volume 32, pages 11918–11930, 2019.
- [32] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- [33] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT Press, Cambridge, MA, 1998.
- [34] Wenpin Tang and Hanyang Zhao. Score-based diffusion models via stochastic differential equations. *Statistics Surveys*, 19:28–64, 2025.

- [35] Wenpin Tang and Xun Yu Zhou. Regret of exploratory policy improvement and q -learning. *arXiv preprint arXiv:2411.01302*, 2024.
- [36] Haoran Wang, Thaleia Zariphopoulou, and Xun Yu Zhou. Reinforcement learning in continuous time and space: A stochastic control approach. *Journal of Machine Learning Research*, 21 (198):1–34, 2020.
- [37] Yuchen Wu, Yuxin Chen, and Yuting Wei. Stochastic Runge-Kutta methods: Provable acceleration of diffusion models. *arXiv preprint arXiv:2410.04760*, 2024.
- [38] Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. In *International Conference on Learning Representations*, 2023.
- [39] Qinsheng Zhang, Jiaming Song, and Yongxin Chen. Improved order analysis and design of exponential integrator for diffusion models sampling. *arXiv preprint arXiv:2308.02157*, 2023.
- [40] Hanyang Zhao, Wenpin Tang, and David D Yao. Policy optimization for continuous reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 36, 2023.
- [41] Hanyang Zhao, Haoxian Chen, Ji Zhang, David D Yao, and Wenpin Tang. Scores as Actions: a framework of fine-tuning diffusion models by continuous-time reinforcement learning. *arXiv preprint arXiv:2409.08400*, 2024.
- [42] Hanyang Zhao, Haoxian Chen, Ji Zhang, David Yao, and Wenpin Tang. Score as Action: Fine tuning diffusion generative models by continuous-time reinforcement learning. In *International Conference on Machine Learning*, 2025.
- [43] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *AAAI Conference on Artificial Intelligence*, volume 8, pages 1433–1438, Chicago, IL, USA, 2008.