

# Hierarchical Schedule Optimization for Fast and Robust Diffusion Model Sampling

Aihua Zhu<sup>1\*</sup>, Rui Su<sup>1\*</sup>, Qinglin Zhao<sup>1†</sup>, Li Feng<sup>1</sup>, Meng Shen<sup>2</sup>, Shibo He<sup>3</sup>

<sup>1</sup>School of Computer Science and Engineering, Macau University of Science and Technology, Macau 999078, China

<sup>2</sup>Beijing Institute of Technology

<sup>3</sup>Zhejiang University

## Abstract

Diffusion probabilistic models have set a new standard for generative fidelity but are hindered by a slow iterative sampling process. A powerful training-free strategy to accelerate this process is *Schedule Optimization*, which aims to find an optimal distribution of timesteps for a fixed and small Number of Function Evaluations (NFE) to maximize sample quality. To this end, a successful schedule optimization method must adhere to four core principles: effectiveness, adaptivity, practical robustness, and computational efficiency. However, existing paradigms struggle to satisfy these principles simultaneously, motivating the need for a more advanced solution. To overcome these limitations, we propose the Hierarchical-Schedule-Optimizer (HSO), a novel and efficient bi-level optimization framework. HSO reframes the search for a globally optimal schedule into a more tractable problem by iteratively alternating between two synergistic levels: an upper-level global search for an optimal initialization strategy and a lower-level local optimization for schedule refinement. This process is guided by two key innovations: the Midpoint Error Proxy (MEP), a solver-agnostic and numerically stable objective for effective local optimization, and the Spacing-Penalized Fitness (SPF) function, which ensures practical robustness by penalizing pathologically close timesteps. Extensive experiments show that HSO sets a new state-of-the-art for training-free sampling in the extremely low-NFE regime. For instance, with an NFE of just 5, HSO achieves a remarkable FID of 11.94 on LAION-Aesthetics with Stable Diffusion v2.1. Crucially, this level of performance is attained not through costly retraining, but with a one-time optimization cost of *less than 8 seconds*, presenting a highly practical and efficient paradigm for diffusion model acceleration.

**Code** — <https://github.com/chappy0/HSO.git>

**Extended version** — <https://arxiv.org/abs/2511.11688>

## 1 Introduction

Diffusion probabilistic models (Ho, Jain, and Abbeel 2020; Song et al. 2021) have established a new standard for high-quality image generation. However, their practical utility is hindered by a computationally intensive iterative sampling

\*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Principle	Rule-based	Perceptual Opt.	Principled Opt.	HSO (Ours)
Adaptivity	✗	✓	✓	✓
Effectiveness	✗	✓	✗	✓
Practical Robustness	✓	✗	✗	✓
Comp. Efficiency	✓	✗	✓	✓

Table 1: A comparison of Schedule Optimization paradigms on four key principles. HSO uniquely satisfies all four.

process. At the heart of this process from pure noise to a clean image, a numerical solver repeatedly invokes the model’s neural network function over a sequence of discrete timesteps, known as the *schedule*. The total number of these invocations termed the Number of Function Evaluations (NFE), often reaches the hundreds or thousands, creating a significant bottleneck for real-time applications. Consequently, this makes the reduction of NFE while maintaining high sample fidelity a primary research goal.

A powerful strategy to address this goal is *Schedule Optimization*. This approach focuses on determining the optimal distribution of a limited set of sampling timesteps across the full generation process from noise to image, which in turn defines the sampling schedule. A key advantage is its **training-free** nature, which refines the sampling process to accelerate diffusion models without costly retraining or fine-tuning. The ultimate objective of this task is to discover a schedule that yields the highest possible sample fidelity with a minimal search cost. To systematically design an approach capable of meeting this objective, we argue that it must adhere to four fundamental principles:

- **Adaptivity:** The process must derive a schedule adapted to the intrinsic characteristics of models and given NFE budgets, rather than relying on a one-size-fits-all rule independent to the models and the NFE.
- **Effectiveness:** The underlying logic of the method must be capable of producing a schedule that results in high-fidelity samples.
- **Practical Robustness:** The approach must incorporate explicit constraints to ensure robustness against the devi-

ation between its theoretical objective and actual fidelity.

- **Computational Efficiency:** The total cost of the search process must be practically viable and acceptable.

However, as summarized in Table 1, current paradigms for schedule optimization struggle to satisfy these four principles simultaneously, each presenting its own trade-offs. **Rule-based methods**, while computationally efficient and often robust, inherently lack Adaptivity as their fixed rules cannot adjust to diverse models and scenarios. On the other hand, optimization-based approaches are adaptive but struggle with the other principles. **Perceptual Optimization**, which directly optimizes for image perceptual quality metrics, sacrifices Computational Efficiency due to the high cost of repeated image generation during the search. **Principled Optimization**, which *seeks to minimize a theoretically-derived objective* instead, are undermined by a non-convex landscape that traps local searches in suboptimal minima, crippling their Effectiveness and often neglecting Practical Robustness.

To resolve these shortcomings and holistically satisfy all aforementioned principles, we introduce the **Hierarchical-Schedule-Optimizer (HSO)**, a novel and efficient bi-level optimization framework designed to navigate the non-convex search landscape. HSO iteratively alternates between an upper-level global search within a low-dimensional space for an optimal initialization strategy and a lower-level local optimization for schedule refinement from that superior initial point. This process is guided by two key innovations: the **Midpoint Error Proxy (MEP)**, a solver-agnostic and numerically stable objective for effective local optimization, and the **Spacing-Penalized Fitness (SPF)** function, which ensures practical robustness by penalizing pathologically close timesteps.

Our main contributions can be summarized as follows:

- We propose HSO, a novel bi-level optimization framework that significantly advances the Principled Optimization paradigm, with a search process that adheres to the principles of Adaptivity, Effectiveness, Practical Robustness, and Computational Efficiency.
- Within the HSO framework, we present two key technical innovations: the MEP guiding the lower-level local optimization, and the SPF function steering the upper-level global search.
- We demonstrate through extensive experiments that our approach sets a new state-of-the-art for training-free sampling in the extremely low-NFE regime. For instance, at an NFE of 5, using the Stable Diffusion v2.1 model on the LAION-Aesthetics 6.5+ benchmark, HSO achieves an FID(Heusel et al. 2017) score as low as 11.94, all with a one-time optimization cost of **less than 8 seconds**, presenting a highly practical paradigm for diffusion model acceleration.

## 2 Related Work

Accelerating diffusion model sampling is a critical area of research. To precisely position our contributions, we first overview the broader acceleration landscape before focusing on prior work in Schedule Optimization.

## 2.1 Approaches to Sampling Acceleration

Methods for accelerating diffusion model sampling can be broadly divided by whether they require retraining the diffusion model itself.

**Training-based acceleration** creates or fine-tunes a diffusion model for few-step sampling. Prominent techniques include knowledge distillation (Geng, Pokle, and Kolter 2023; Salimans and Ho 2022; Sauer et al. 2024; Luo et al. 2023b; Yin et al. 2024; Zhou et al. 2024; Frans et al. 2025; Lin, Wang, and Yang 2024; Xu et al. 2024; Watson et al. 2021a,b), consistency models (Song et al. 2023; Geng et al. 2024; Lu and Song 2025; Luo et al. 2023a; Song and Dhariwal 2023; Yang et al. 2024), and the advanced work Mean Flows (Geng et al. 2025), which can achieve extremely low NFE, even single-step generation. Their primary drawback, however, is the substantial computational cost required for this training or fine-tuning process, which limits their practical applicability to new or existing models. **Training-free acceleration**, aims to speed up any pre-trained model without modifying its weights. This area encompasses: 1) Solver Design, which focuses on creating custom ODE/SDE solvers or using advanced test-time guidance (Zhao et al. 2023; Lu et al. 2022, 2025; Frankel et al. 2025; Gonzalez et al. 2023; Wang et al. 2025; Xia et al. 2023; Kim, Kim, and Park 2025); 2) Model-level, which optimizes the neural network’s architecture through techniques like quantization or caching (Shang et al. 2023; Liu, Li, and Gu 2025); 3) Schedule Optimization, the focus of our work.

## 2.2 Schedule Optimization Strategies

The task of schedule optimization is to select an optimal subsequence of timesteps. Existing strategies are typically either rule-based or optimization-based.

**Rule-based schedules** provide a fixed, pre-defined formula for generating timesteps, such as the power-law distribution used in the EDM schedule (Karras et al. 2022) and the uniform log-SNR schedule defaulted by DPM-Solver (Lu et al. 2022). **Optimization-based schedules** formulate the task as a search problem, aiming to find a schedule that optimizes a specific objective function. This category can be further divided into two main branches: **perceptual optimization** methods directly optimize for image quality metrics such as FID or KID (Li et al. 2023; Liu et al. 2023; Wang et al. 2023); **principled optimization** method like DM-NonUni (Xue et al. 2024) and Score-Optimal Diffusion Schedules (Williams et al. 2024), which is most similar to our work, aim to minimize a theoretically-derived objective.

## 3 Preliminaries: Foundations of Principled Schedule Optimization

This section establishes the technical foundations for Principled Schedule Optimization by defining the schedule optimization problem within the probability flow ODE framework, introducing a parameterized method for constructing an initial schedule, and outlining the principle of a theoretical optimization objective.

### 3.1 The Schedule Optimization Problem

A diffusion model’s forward process transforms a clean data sample  $\mathbf{x}_0$  into pure noise over a continuous time  $t \in [0, T]$ . The distribution of the noisy sample  $\mathbf{x}_t$  conditioned on  $\mathbf{x}_0$  is a Gaussian:  $q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\alpha_t\mathbf{x}_0, \sigma_t^2\mathbf{I})$ , where  $\alpha_t$  and  $\sigma_t$  are the signal and noise scales, respectively.

The reverse generation process can be described by a probability flow Ordinary Differential Equation (ODE) for deterministic sampling (Song et al. 2021). For a model with a data prediction network  $\mathbf{x}_\theta(\mathbf{x}_t, t)$  to approximate the original clean data  $\mathbf{x}_0$ , this ODE takes the specific form (Xue et al. 2024):

$$\frac{d\mathbf{x}_t}{dt} = \frac{d \log \sigma_t}{dt} \mathbf{x}_t + \frac{\alpha_t d \log \frac{\alpha_t}{\sigma_t}}{dt} \mathbf{x}_\theta(\mathbf{x}_t, t) \quad (1)$$

A numerical ODE solver approximates the solution to this equation by evaluating  $\mathbf{x}_\theta$  at a discrete sequence of timesteps  $\tau = \{t_0, t_1, \dots, t_N\}$ , where  $t_0 = T$  is the start time and  $t_N = \epsilon$  is the end time for some small  $\epsilon > 0$ . For improved numerical stability, this process is often performed in the log-signal-to-noise ratio (log-SNR) space, defined by  $\lambda_t = \log(\alpha_t/\sigma_t)$ . The core problem of schedule optimization is therefore to find the optimal discrete set of log-SNR values,  $\Lambda = \{\lambda_0, \lambda_1, \dots, \lambda_N\}$ , that minimizes the final generation error for a fixed NFE budget  $N$ .

### 3.2 Parameterizing the Optimization Initial Point

Optimization-based method requires an initial schedule  $\Lambda_{\text{init}}$ , or a initial point, from which to begin its search. A structured way to construct this schedule is by generating it from a low-dimensional hyperparameter vector, denoted by  $\psi$ . The highly effective EDM schedule proposed by Karras et al. (2022) provides an excellent formula for this purpose. Their method defines the schedule based on a noise level parameter,  $\tilde{\sigma}$ , which relates to the total noise and signal scales via  $\sigma_t = \alpha_t \tilde{\sigma}_t$ . The schedule for these discrete noise level parameters  $\{\tilde{\sigma}_i\}_{i=0}^N$  is constructed by the formula ( $\tilde{\sigma}_N = 0$ ):

$$\tilde{\sigma}_i = \left( \tilde{\sigma}_{\min}^{1/\rho} + \frac{i}{N-1} \left( \tilde{\sigma}_{\max}^{1/\rho} - \tilde{\sigma}_{\min}^{1/\rho} \right) \right)^\rho \quad (2)$$

This entire construction is governed by the compact hyperparameter vector  $\psi = (\rho, \tilde{\sigma}_{\min}, \tilde{\sigma}_{\max}) \in \mathbb{R}^3$ , which defines the schedule’s distribution and boundaries. By converting these  $\tilde{\sigma}_i$  values to the log-SNR space via  $\lambda_i = -\log(\tilde{\sigma}_i)$ , we can generate a complete log-SNR schedule, denoted as  $\Lambda(\psi)$ . This parameterized function allows a complete initial schedule,  $\Lambda_{\text{init}} = \Lambda(\psi)$ , to be generated from a compact set of inputs. The low-dimensional vector  $\psi$  thus represents the “initialization strategy” that HSO’s upper level is designed to optimize.

### 3.3 Principle of a Theoretical Optimization Objective

Principled schedule optimization methods require a theoretical objective function,  $\mathcal{J}_{\text{theo}}(\Lambda)$ , to guide the search for an optimal schedule. The core idea, pioneered by DM-NonUni, is to derive this objective from the global generation error  $\|\tilde{x}_\epsilon - x_0\|$ , where  $\tilde{x}_\epsilon$  is the numerical approximation of the

true solution  $x_\epsilon$ . This error originates from the integral solution of the ODE in equation 1:

$$x_\epsilon = \frac{\sigma_\epsilon}{\sigma_T} x_T + \sigma_\epsilon \sum_{i=0}^{N-1} \int_{\lambda_i}^{\lambda_{i+1}} e^\lambda f(\lambda) d\lambda \quad (3)$$

where  $f(\lambda) = \mathbf{x}_\theta(\mathbf{x}_{t(\lambda)}, t(\lambda))$ . A tractable objective can be formulated by first approximating the integral term. The specific form of this objective depends on the chosen approximation method for  $f(\lambda)$  within the integral. This principle provides the theoretical foundation for the novel objective function we introduce in our Section 4.2.

## 4 Methodology: The HSO

This section details the Hierarchical-Schedule-Optimizer (HSO), our framework designed to systematically satisfy the four foundational principles of schedule optimization. We will first present its general bi-level architecture, then detail its two core components, the Midpoint Error Proxy (MEP) and the Spacing-Penalized Fitness (SPF), before concluding with the complete algorithm.

### 4.1 The HSO Framework: A Bi-Level Solution to a Non-Convex Problem

The goal of principled schedule optimization is to find a schedule  $\Lambda$  that minimizes a theoretical objective function  $\mathcal{J}_{\text{theo}}(\Lambda)$  derived from the global generation error in Section 3.3. However, the optimization landscape  $\mathcal{J}_{\text{theo}}(\Lambda)$  is highly non-convex, making local optimization methods highly sensitive to their initial point  $\Lambda_{\text{init}}$  and prone to getting trapped in suboptimal minima (Jain, Kar et al. 2017).

The direct consequence of this non-convex landscape is that the success of local optimization becomes critically dependent on finding a superior initial point. While a global search is the most robust method for locating such a point, searching for a schedule  $\Lambda_{\text{init}} \in \mathbb{R}^N$  directly in the  $N$ -dimensional schedule space is often computationally intractable due to the curse of dimensionality, especially as the NFE  $N$  increases. To overcome this tractability barrier, HSO reframes the problem from a direct, high-dimensional search over  $\Lambda$  to a more tractable, global search for a single point  $\psi$  that generates the initial schedule within the low-dimensional hyperparameter space  $\mathbb{R}^3$ . This reframing naturally gives rise to the HSO bi-level framework, which decomposes the difficult optimization for a global optimal schedule into two manageable sub-problems: a global search for an optimal initialization strategy and a local optimization of the resulting schedule, which are solved through an iterative, alternating process between the two levels of our framework:

- **Lower-Level Local Optimization:** This level performs a local search for an optimal schedule  $\Lambda_{\text{opt}}$ , starting from an initial point  $\Lambda_{\text{init}}$  which is generated using the hyperparameter vector  $\psi$  from the upper level. This search is guided by a flexible theoretical objective function  $\mathcal{J}_{\text{lower}}$  and executed using a standard constrained optimization algorithm, such as the Trust-Region Constrained Opti-

mization method (Conn, Gould, and Toint 2000). Formally, its task is to solve:

$$\Lambda_{\text{opt}}(\psi) = \arg \min_{\Lambda} \mathcal{J}_{\text{lower}}(\Lambda | \Lambda_{\text{init}}(\psi)) \quad (4)$$

- **The Upper-level Global Search:** This level performs a global search for the optimal hyperparameter vector  $\psi^*$  that defines the best initial point. It employs a population-based evolutionary algorithm, such as Differential Evolution (Storn and Price 1997), to effectively navigate the non-convex landscape. This search is guided by a flexible, practicality-aware fitness function  $\mathcal{F}_{\text{upper}}$ , which evaluates the quality of the final schedule  $\Lambda_{\text{opt}}(\psi)$  produced by the lower level for a given  $\psi$ . The task can be formally expressed as:

$$\psi^* = \arg \min_{\psi} \mathcal{F}_{\text{upper}}(\Lambda_{\text{opt}}(\psi)) \quad (5)$$

This bi-level optimization is realized through an iterative, alternating process. In each cycle, the upper level proposes a population of candidate strategies  $\{\psi_j\}$ . For each candidate, the lower level executes a local optimization to yield a corresponding schedule  $\{\Lambda_{\text{opt},j}\}$ . The quality of these schedules is then evaluated and fed back to the upper level to guide the evolution of its population for the next cycle. This iterative process of global search and local optimization continues until a termination condition is met.

The performance of the HSO framework is ultimately driven by the specific designs of its two guiding functions, the lower-level objective  $\mathcal{J}_{\text{lower}}$  and the upper-level fitness function  $\mathcal{F}_{\text{upper}}$ , which we detail in the following subsections.

#### 4.2 Objective for Local Optimization: The Midpoint Error Proxy (MEP)

The design of the lower level’s theoretical objective  $\mathcal{J}_{\text{lower}}$ , critically influences the Effectiveness, Adaptivity and Computational Efficiency of the HSO framework. While DM-NonUni derives objectives tightly coupled to the mechanics of specific high-order solvers like UniPC (Zhao et al. 2023), this approach can limit generality and numerical stability, as the underlying extrapolation-based polynomial integral approximation used in UniPC sampling can produce large estimation errors and lead to instability (Burden and Faires 1993). As a core theoretical contribution of this work, we diverge from this solver-specific path and propose a more universal and numerically reliable objective, the Midpoint Error Proxy (MEP)  $\mathcal{J}_{\text{MEP}}$ , to serve as the lower level’s objective  $\mathcal{J}_{\text{lower}}$ .

Following the principle from Section 3.3, our goal is to derive a tractable objective by approximating the integral term in the global error formula shown in Eq.(3). Our approach is use an interpolation-based approximation inspired by the classic midpoint rule, based on the principle of isolating and exactly integrating the analytically tractable exponential term  $e^\lambda$ , while only applying approximation to the neural network term  $f(\lambda)$  which we formalize in Lemma 1.

**Lemma 1** (Hybrid Midpoint Approximation). *The integral term can be approximated as:*

$$\int_{\lambda_i}^{\lambda_{i+1}} e^\lambda f(\lambda) d\lambda \approx f(\lambda_{i+\frac{1}{2}})(e^{\lambda_{i+1}} - e^{\lambda_i}) \quad (6)$$

where  $\lambda_{i+\frac{1}{2}} \triangleq \frac{\lambda_i + \lambda_{i+1}}{2}$ . *The local truncation error of this approximation is  $O(h^3)$ , the same order as the standard midpoint rule, but it is numerically more stable for this specific integral structure.*

A detailed derivation and proof are provided in the Appendix A. Building on this lemma, we can now establish the global error bound. To formalize this bound, we adopt the theoretical setup from DM-NonUni, assuming the model’s prediction is  $f(\lambda) = x_0 + \xi_{t(\lambda)}$ , where the error norm is bounded with high probability by  $\|\xi_{t(\lambda)}\| \leq \tilde{\eta} \tilde{\epsilon}_{t(\lambda)}$  for some constant  $\tilde{\eta}$  and error proxy  $\tilde{\epsilon}_{t(\lambda)} = \frac{\sigma_t^p}{\alpha_t}$  for a non-negative integer  $p$ .

**Theorem 1** (Global Error Bound for Hybrid Midpoint Approximation). *Under the aforementioned assumptions, let the approximate solution  $\tilde{x}_{\epsilon, \text{MEP}}$  be derived using the approximation from Lemma 1. The global generation error is then bounded with high probability by:*

$$\|\tilde{x}_{\epsilon, \text{MEP}} - x_0\| \leq C + \sigma_\epsilon \tilde{\eta} \sum_{i=0}^{N-1} \tilde{\epsilon}_{t(\lambda_{i+\frac{1}{2}})} (e^{\lambda_{i+1}} - e^{\lambda_i}) \quad (7)$$

where  $C$  represents schedule-independent error terms.

A proof is provided in the Appendix B. Based on Theorem 1, minimizing the error bound is achieved by minimizing the schedule-dependent summation. This directly yields our MEP objective:

$$\mathcal{J}_{\text{MEP}}(\Lambda) = \sum_{i=0}^{N-1} \tilde{\epsilon}_{t(\lambda_{i+\frac{1}{2}})} (e^{\lambda_{i+1}} - e^{\lambda_i}) \quad (8)$$

**Remark 1** (Adherence to Core Principles). *The MEP objective is deliberately designed to embody our core principles. Its high-order accuracy and numerical reliability ensure **Effectiveness**. It ensures **Adaptivity** because the error proxy  $\tilde{\epsilon}_{t(\lambda)}$  incorporates the model’s unique noise schedule parameters and the optimization is performed for a specific NFE. Finally, its clean analytical form is computable in linear time  $O(N)$ , ensures **Computational Efficiency**.*

#### 4.3 Fitness Function For Global Search: The Spacing-Penalized Fitness (SPF)

The design of the upper level’s fitness function  $\mathcal{F}_{\text{upper}}$  is crucial for ensuring the Practical Robustness of the discovered schedules. Our empirical analysis reveals that unconstrained optimization can produce schedules with pathologically close timesteps. While theoretically optimal, these schedules are practically ineffective, as the close steps provide negligible progress and waste the NFE budget.

To address this, we design the Spacing-Penalized Fitness (SPF)  $\mathcal{F}_{\text{SPF}}$  to serve as  $\mathcal{F}_{\text{upper}}$ . It augments the theoretical error from our MEP objective with a dynamic spacing penalty,  $L_{\text{penalty}}$ . The combined fitness function is:

---

**Algorithm 1: The Complete HSO Optimization Process**

---

**Input:** Target NFE  $N$ , Search bounds for strategy vector  $\psi$ **Output:** Optimized schedule  $\Lambda^*$ 

```
1: Initialize a set of candidate hyperparameter vectors  $P =$   
    $\{\psi_1, \dots, \psi_K\}$   $\triangleright$  Init each  $\psi_k \in \mathbb{R}^3$   
2: Initialize best fitness  $f^* \leftarrow \infty$ , best schedule  $\Lambda^* \leftarrow \text{null}$   
3: while TermCond is not met do  $\triangleright$  // Global Search  
4:   for each candidate  $\psi_k$  in  $P$  do  $\triangleright$  // Local Opt.  
5:      $\Lambda_{\text{init}} \leftarrow \text{GenerateInitialPoint}(\psi_k)$   
6:      $\Lambda_{\text{opt},k} \leftarrow \text{Minimize}(\mathcal{J}_{MEP}, \Lambda_{\text{init}})$   
        $\triangleright$  // Minimize  $\mathcal{J}_{MEP}$  with  $\Lambda_{\text{init}}$  as start  
7:      $f_k \leftarrow \mathcal{F}_{\text{SPF}}(\Lambda_{\text{opt},k})$   $\triangleright$  // Evaluate fitness of  $\Lambda_{\text{opt},k}$   
8:     if  $f_j < f^*$  then  
9:        $f^* \leftarrow f_k$   
10:       $\Lambda^* \leftarrow \Lambda_{\text{opt},k}$   
11:     end if  
12:   end for  
13:    $P \leftarrow \text{EvolvePopulation}(P, \{f_1, \dots, f_K\})$   
14: end while
```

---

$$\mathcal{F}_{\text{SPF}}(\psi|N) = \mathcal{J}_{\text{MEP}}(\Lambda_{\text{opt}}(\psi)) + \gamma L_{\text{penalty}}(\Lambda_{\text{opt}}(\psi)|N) \quad (9)$$

where the penalty term heavily penalizes schedules when timesteps cluster closer than a minimum distance  $d_{\min}(N)$  that adapts based on the NFE budget:

$$L_{\text{penalty}} = \sum_{i=0}^{N-1} \max(0, d_{\min}(N) - |t(\lambda_{i+1}) - t(\lambda_i)|)^2 \quad (10)$$

The value of  $d_{\min}(N)$  adapts to the NFE budget based on a simple linear heuristic, decreasing from 0.15 at NFE  $N = 4$  down to 0.01 at  $N = 20$ . This strategy enforces a wider step spacing to ensure robustness at very low NFEs, while permitting finer-grained steps for higher precision when a larger evaluation budget is available.

## 4.4 HSO Implementation and Algorithm

We summarize the complete HSO algorithm, built on our core components MEP and SPF, in Algorithm 1. The synergy between its efficient, low-dimensional global search and robust, MEP-guided local optimization objective regularized by the SPF’s spacing penalty, ensures the discovery of schedules that are simultaneously Effective, Adaptive, Practically Robust, and Computationally Efficient.

## 5 Experiments

This section experimentally validates our Hierarchical Schedule Optimizer (HSO). Following the setup, we demonstrate its Effectiveness, Adaptivity, Practical Robustness, and Computational Efficiency. Finally, an ablation study verifies the individual and synergistic contributions of HSO’s core components.

### 5.1 Experimental Setup

For reproducibility, we evaluate HSO on the Stable Diffusion v2.1-base model, a widely-adopted open-source model,

using two prominent ODE solvers, UniPC (Zhao et al. 2023) and DDIM (Song, Meng, and Ermon 2020), assessing its generalization capabilities across three standard, large-scale benchmarks: LAION-Aesthetics 6.5+ (comprising approximately 30,000 text-image pairs)(Schuhmann et al. 2022), the MS-COCO 2017 validation set (with 30,000 captions)(Lin et al. 2014), and the ImageNet 512x512 validation set (with 50,000 captions)(Deng et al. 2009). We measure image quality using the Fréchet Inception Distance (FID) (Heusel et al. 2017), a standard metric for perceptual fidelity. The main parameters for HSO, especially the search boundaries for the strategy vector  $\psi$  and our rationale for this, are detailed in the Appendix C. We benchmark HSO against SOTA acceleration methods, especially the principled optimization DM-NonUni(Xue et al. 2024) focusing on the demanding low-NFE ( $N < 5$ ) regime.

### 5.2 Validation of the Effectiveness

**Quantitative analysis.** To validate its effectiveness, we use FID(Heusel et al. 2017) to quantitatively measure sample fidelity, where a lower score indicates higher quality. As shown in Table 2, HSO consistently and significantly outperforms the baseline DM-NonUni across all configurations. This advantage is most pronounced in the extremely low-NFE regime. For instance, on LAION-Aesthetics with the DDIM solver, as NFE drops from 5 to 4, the baseline FID collapses from 35.38 to 68.92, whereas HSO’s degrades more gracefully from 17.17 to 24.77. While HSO shows clear improvements on UniPC, the performance gain is substantially larger on the DDIM solver (e.g., on LAION, FID drops from 68.92 to 24.77 at  $NFE = 4$ ). This is because the baseline’s objective was tailored for UniPC, whereas our solver-agnostic MEP objective generalizes more effectively, thereby validating the universality of our approach. This robust performance across different solvers, especially in the demanding low-NFE regime, confirms HSO’s ability to find highly effective schedules.

Dataset	Solver	NFE	FID↓ (CVPR)	FID↓ HSO (Ours)
LAION-Aesthetics 6.5+	UniPC	4	18.96	<b>15.71</b>
		5	13.91	<b>11.94</b>
	DDIM	4	68.92	24.77
		5	35.38	17.17
MS-COCO	UniPC	4	27.50	23.26
		5	23.10	19.59
	DDIM	4	60.06	29.55
		5	30.12	23.15
ImageNet (512×512)	UniPC	4	20.75	17.20
		5	17.63	15.90
	DDIM	4	41.51	19.78
		5	22.89	16.02

Table 2: Performance comparison (FID↓) between HSO and DM-NonUni. The “FID (CVPR)” column refers to the results with the DM-NonUni method (Xue et al. 2024).

**Qualitative Analysis.** The qualitative results in Figure 1 visually corroborate our quantitative advantages over base-

line DM-NonUni. HSO consistently generates images with high semantic accuracy and detail fidelity across different datasets and solvers. For instance, it accurately renders the “astronaut cat” (Fig. 1a), the “vintage car on a coast” (Fig. 1b), and a realistic “golden retriever” (Fig. 1c). In contrast, the baseline method struggles at low NFEs, leading to both severe artifacts and semantic failures. This is evident in the distorted cat-helmet fusion (Fig. 1d), the failure to generate the correct coastal background for the car (Fig. 1e), and the malformed retriever face (Fig. 1f).



Figure 1: Visual comparison of HSO (top) vs. DM-NonUni (bottom). HSO demonstrates superior image quality and coherence, especially at extremely low NFEs.

### 5.3 Validation of Adaptivity

HSO demonstrates strong adaptivity to both computational budgets (NFE) and model noise schedules.

**NFE Adaptivity.** We hypothesize the optimization landscape is highly non-convex, making any fixed rule-based approach unlikely to be optimal across all scenarios. The optimization-based nature for HSO made it able to tailor the strategy adaptively according to the given. To validate this, we tasked HSO to find the optimal parameters  $\psi^*$  for various NFE values. The results, presented in Table 3, strikingly demonstrate this adaptive behavior. The optimal parameters exhibit a complex, non-monotonic relationship with NFE (e.g.,  $\rho^*$  peaks at NFE=8), refuting any “one-size-fits-all” approach and highlighting the necessity of HSO’s adaptive search.

**Model Adaptivity.** HSO’s adaptivity also extends to the unique noise characteristics inherent in different models. Theoretically, any two models with distinct noise progressions (i.e., different mappings from timestep  $t$  to noise levels  $\alpha_t$  and  $\sigma_t$ ) require different optimal sampling schedules, even if they belong to the same model family. To validate this, we test HSO on PixArt- $\alpha$  on the MS-COCO

NFE	$t_{\min}^*$	$t_{\max}^*$	$\rho^*$	FID
4	0.03	0.96	8.8431	15.71
6	0.010 001 87	0.999 936 8	6.5085	11.17
8	0.010 000 35	0.999 894 1	12.4163	8.79
10	0.010 251 1	0.999 952 7	11.6272	8.32

Table 3: Optimal initialization parameters discovered by HSO for varying NFEs. We search for  $t$  instead of  $\sigma$ , as they share a one-to-one correspondence (Appendix C).

30k validation set, using the UniPC solver with an NFE of 5. Although both PixArt- $\alpha$  and Stable Diffusion v2.1 are Variance-Preserving (VP) models, they employ different  $\beta$  schedules, resulting in distinct model properties that HSO can adapt to. As shown in Table 4, applying HSO to PixArt- $\alpha$  significantly improves the FID score over its default schedule. This result demonstrates that HSO effectively optimizes for the specific characteristics of a given model, underscoring its versatility.

Foundational Model	Baseline FID	HSO FID
PixArt- $\alpha$ (512px)	37.65	18.05
SD-V2.1 (512px)	24.06	19.59

Table 4: HSO performance on PixArt- $\alpha$  (512px), evaluated with NFE=5 and the UniPC solver on the MS-COCO 30k validation set.

### 5.4 Validation of Practical Robustness

In the following experiments with the DDIM scheduler, we use the standard discretization where the continuous time range is mapped to 1000 integer steps (indexed 0-999). Unconstrained optimization can yield “brittle” schedules with pathologically close timesteps that hurt empirical performance. For instance, without our SPF, the DDIM schedule collapsed to [999, 70, 9, 9], exhibiting an extreme “end-clustering” of steps, resulting in a catastrophic FID of 165.48 (Table 5). The minimum spacing between steps collapsed to 0.0, a clear indicator of numerical instability, which resulted in an unusable FID of **165.48**. In contrast, enabling our NFE-adaptive SPF penalty enforces a healthy minimum step spacing (243.0) and reliably guides the optimizer toward high-performing, stable solutions, achieving a stable mean FID of **19.76±0.25** over 10 runs. This confirms our regularization is crucial for practical robustness.

### 5.5 Validation of the Computational Efficiency

A core design principle of HSO is practical computational efficiency. This section validates HSO’s **preparation cost** against other leading “Training-Free” and “Training-Based” methods, as shown in Table 6. HSO’s efficiency is a key advantage, finding its schedule in  $\sim 8$  seconds on a consumer-grade CPU. Within the training-free landscape, while slightly longer than the purely local search of DM-NonUni ( $\sim 1$ s), this is an expected trade-off for our bi-level

Condition	Example Schedule	Step Spacing (Min / Mean)	FID
Without Penalty	[999, 70, 9, 9]	0.0 / 330.0	<b>165.48</b> (Unstable)
With Penalty	[959, 716, 370, 30]	243.0 / 309.7	<b>19.76 ± 0.25</b> (~10 runs)

Table 5: Comparison of schedules optimized with and without a penalty. The penalty-based regularization prevents step collapse and leads to stable, high-quality results.

global search and is orders of magnitude faster than other global methods like AutoDiffusion (~1.1 days). Crucially, this low cost is paired with a leading FID score on the challenging ImageNet 512x512 benchmark, whereas competitors often report on lower-resolution 256x256 images. As a training-free method, HSO’s 8-second cost bypasses the substantial training investments (often GPU-days or months) required by training-based paradigms.

	Method	Paradigm	Prep. Cost <sup>‡</sup>	NFE	FID
<i>Training-Free Methods</i>	<b>HSO (Ours)</b>	<b>Hierarchical Opt.</b>	~8s	<b>4</b>	15.71 <sup>†</sup>
	AutoDiffusion [ICCV ’23]	Evolutionary Search	~1.1 d	4	17.86
	DM-NonUni [CVPR ’24]	Local Optimization	~1s	4	18.96
<i>Training-Based Methods</i>	UFOGen [CVPR ’24]	Diffusion-GAN	~12 d	1	22.50 <sup>*</sup>
	Mean Flows [ICLR ’25]	Flow Matching	~60 d	1	3.43
	LCM [ICCV ’23]	Distillation	~1.33 d	4	11.10 <sup>§</sup>

<sup>‡</sup> Units are seconds (s) and days (d) to emphasize the magnitude of difference in cost. FID on ImageNet 256×256 unless noted; <sup>†</sup>ImageNet 512×512; <sup>\*</sup>MSCOCO; <sup>§</sup>LAION-Aesthetics (SD v2.1).

Table 6: Performance comparison with SOTA methods. Prep. cost is estimated in A100 GPU-days (d), except for HSO and DM-NonUni, which were measured on an Intel Core i7-14650HX.

## 5.6 Ablation Study

We conducted an ablation study on the LAION-Aesthetics 6.5+ dataset with NFE=4 to isolate the contributions of HSO’s key components: the Bi-Level search and the MEP objective. All experiments here include the SPF penalty, with its ablation shown in Section 5.4. Results are detailed in Table 7.

**The baseline (A):** The principled local schedule optimization approach proposed by Xue et al.(2024), which is an objective function designed specifically for solvers like Unipc.

**A+Bi-level (B):** Adding only our Bi-Level global search finds a superior initial schedule. This dramatically improves FID for both solvers, especially for DDIM (71.57 → 29.22), as its default performance is far from optimal, offering more room for improvement from a better initialization compared to the already well-performing UniPC.

**A + MEP Objective (C):** This configuration replaces the baseline’s UniPC-specific local objective with our solver-agnostic MEP. This reveals a predictable trade-off: the acceptable performance decrease on UniPC (18.07 → 26.33) is outweighed by a massive improvement on DDIM (71.57 → 37.27). This result strongly validates MEP’s effectiveness as a universal error proxy that avoids overfitting to a single solver.

**Synergy in the Full HSO Framework (Ours):** The complete HSO framework demonstrates powerful synergy. The Bi-Level stage provides a superior foundation, which the MEP objective then refines to achieve the best overall performance. While the UniPC FID is slightly higher than in configuration (B) due to the objective trade-off described in (C), the final DDIM FID reaches a new best of 24.80. This proves our hierarchical design is essential for discovering robust, high-performing schedules.

Component / Configuration	UniPC FID	DDIM FID
Baseline (A)	18.07	71.57
A + Bi-Level (B)	11.44	29.22
A + MEP Objective (C)	26.33	37.27
Synergistic Framework (Ours HSO)	15.70	<b>24.80</b>

Table 7: Ablation study of HSO’s components on UniPC and DDIM solvers (NFE=4).

## 6 Conclusion

In this paper, we addressed the challenge of accelerating diffusion models via training-free schedule optimization. We proposed the Hierarchical-Schedule-Optimizer (HSO), a novel bi-level framework designed to holistically satisfy the principles of Adaptivity, Effectiveness, Practical Robustness, and Computational Efficiency. HSO navigates the non-convex search landscape by synergizing a low-dimensional global search with a rapid local refinement, a process powered by our two innovations: the solver-agnostic Midpoint Error Proxy (MEP) objective and the stability-enhancing Spacing-Penalized Fitness (SPF) function. Extensive experiments validate that HSO establishes a new state-of-the-art in the low-NFE regime with a negligible optimization cost of just a few seconds, confirming it as a practical and powerful solution. For future work, the HSO framework could be extended to co-optimize other sampling components, such as solver hyperparameters, in a unified manner.

## Acknowledgments

This work is funded by the Science and Technology Development Fund, Macau SAR (File No. 0008/2025/RIB1, 0126/2025/RIA2, and 0077/2025/RIA2).

## References

- Burden, R. L.; and Faires, D. J. 1993. *Numerical Analysis*. Boston, MA: PWS, fifth edition.
- Conn, A. R.; Gould, N. I.; and Toint, P. L. 2000. *Trust region methods*. SIAM.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Frankel, E.; Chen, S.; Li, J.; Koh, P. W.; Ratliff, L. J.; and Oh, S. 2025. S4s: Solving for a diffusion model solver. *arXiv preprint arXiv:2502.17423*.
- Frans, K.; Hafner, D.; Levine, S.; and Abbeel, P. 2025. One Step Diffusion via Shortcut Models. In *The Thirteenth International Conference on Learning Representations*.
- Geng, Z.; Deng, M.; Bai, X.; Kolter, J. Z.; and He, K. 2025. Mean flows for one-step generative modeling. *arXiv preprint arXiv:2505.13447*.
- Geng, Z.; Pokle, A.; and Kolter, J. Z. 2023. One-step diffusion distillation via deep equilibrium models. *Advances in Neural Information Processing Systems*, 36: 41914–41931.
- Geng, Z.; Pokle, A.; Luo, W.; Lin, J.; and Kolter, J. Z. 2024. Consistency models made easy. *arXiv preprint arXiv:2406.14548*.
- Gonzalez, M.; Fernandez Pinto, N.; Tran, T.; Hajri, H.; Mas-moudi, N.; et al. 2023. Seeds: Exponential sde solvers for fast high-quality sampling from diffusion models. *Advances in Neural Information Processing Systems*, 36: 68061–68120.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Jain, P.; Kar, P.; et al. 2017. Non-convex optimization for machine learning. *Foundations and Trends® in Machine Learning*, 10(3-4): 142–363.
- Karras, T.; Aittala, M.; Aila, T.; and Laine, S. 2022. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35: 26565–26577.
- Kim, S.; Kim, M.; and Park, D. 2025. Test-time alignment of diffusion models without reward over-optimization. *arXiv preprint arXiv:2501.05803*.
- Li, L.; Li, H.; Zheng, X.; Wu, J.; Xiao, X.; Wang, R.; Zheng, M.; Pan, X.; Chao, F.; and Ji, R. 2023. Autodiffusion: Training-free optimization of time steps and architectures for automated diffusion model acceleration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7105–7114.
- Lin, S.; Wang, A.; and Yang, X. 2024. Sdxl-lightning: Progressive adversarial diffusion distillation. *arXiv preprint arXiv:2402.13929*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, E.; Ning, X.; Lin, Z.; Yang, H.; and Wang, Y. 2023. Oms-dpm: Optimizing the model schedule for diffusion probabilistic models. In *International Conference on Machine Learning*, 21915–21936. PMLR.
- Liu, X.; Li, Z.; and Gu, Q. 2025. Cachequant: Comprehensively accelerated diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 23269–23280.
- Lu, C.; and Song, Y. 2025. Simplifying, Stabilizing and Scaling Continuous-time Consistency Models. In *The Thirteenth International Conference on Learning Representations*.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35: 5775–5787.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2025. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *Machine Intelligence Research*, 1–22.
- Luo, S.; Tan, Y.; Huang, L.; Li, J.; and Zhao, H. 2023a. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*.
- Luo, W.; Hu, T.; Zhang, S.; Sun, J.; Li, Z.; and Zhang, Z. 2023b. Diff-instruct: A universal approach for transferring knowledge from pre-trained diffusion models. *Advances in Neural Information Processing Systems*, 36: 76525–76546.
- Salimans, T.; and Ho, J. 2022. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*.
- Sauer, A.; Lorenz, D.; Blattmann, A.; and Rombach, R. 2024. Adversarial diffusion distillation. In *European Conference on Computer Vision*, 87–103. Springer.
- Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in neural information processing systems*, 35: 25278–25294.
- Shang, Y.; Yuan, Z.; Xie, B.; Wu, B.; and Yan, Y. 2023. Post-training quantization on diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1972–1981.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.

- Song, Y.; and Dhariwal, P. 2023. Improved techniques for training consistency models. *arXiv preprint arXiv:2310.14189*.
- Song, Y.; Dhariwal, P.; Chen, M.; and Sutskever, I. 2023. Consistency models.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Storn, R.; and Price, K. 1997. Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization*, 11(4): 341–359.
- Wang, G.; Cai, Y.; Iijiang Li; Peng, W.; and Su, S.-Z. 2025. PFDiff: Training-Free Acceleration of Diffusion Models Combining Past and Future Scores. In *The Thirteenth International Conference on Learning Representations*.
- Wang, Y.; Wang, X.; Dinh, A.-D.; Du, B.; and Xu, C. 2023. Learning to schedule in diffusion probabilistic models. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2478–2488.
- Watson, D.; Chan, W.; Ho, J.; and Norouzi, M. 2021a. Learning fast samplers for diffusion models by differentiating through sample quality. In *International Conference on Learning Representations*.
- Watson, D.; Ho, J.; Norouzi, M.; and Chan, W. 2021b. Learning to efficiently sample from diffusion probabilistic models. *arXiv preprint arXiv:2106.03802*.
- Williams, C.; Campbell, A.; Doucet, A.; and Syed, S. 2024. Score-Optimal Diffusion Schedules. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, 107960–107983. Neural Information Processing Systems Foundation, Inc.
- Xia, M.; Shen, Y.; Lei, C.; Zhou, Y.; Yi, R.; Zhao, D.; Wang, W.; and Liu, Y.-j. 2023. Towards more accurate diffusion model acceleration with a timestep aligner. *arXiv preprint arXiv:2310.09469*.
- Xu, Y.; Zhao, Y.; Xiao, Z.; and Hou, T. 2024. Ufogen: You forward once large scale text-to-image generation via diffusion gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8196–8206.
- Xue, S.; Liu, Z.; Chen, F.; Zhang, S.; Hu, T.; Xie, E.; and Li, Z. 2024. Accelerating diffusion sampling with optimized time steps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8292–8301.
- Yang, L.; Zhang, Z.; Zhang, Z.; Liu, X.; Xu, M.; Zhang, W.; Meng, C.; Ermon, S.; and Cui, B. 2024. Consistency flow matching: Defining straight flows with velocity consistency. *arXiv preprint arXiv:2407.02398*.
- Yin, T.; Gharbi, M.; Zhang, R.; Shechtman, E.; Durand, F.; Freeman, W. T.; and Park, T. 2024. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6613–6623.
- Zhao, W.; Bai, L.; Rao, Y.; Zhou, J.; and Lu, J. 2023. UniPC: A Unified Predictor-Corrector Framework for Fast Sampling of Diffusion Models. *NeurIPS*.
- Zhou, M.; Zheng, H.; Wang, Z.; Yin, M.; and Huang, H. 2024. Score identity distillation: Exponentially fast distillation of pretrained diffusion models for one-step generation. In *Forty-first International Conference on Machine Learning*.