Temporal Concentration from Rollout Errors: Implicit Preference Optimization for Text-to-Video Generation

Henglin Liu[1,2,*], **Fangyuan Kong**[2,§], **Jing Wang**[2,3,*], **Yizhou Lin**[1],
Nisha Huang[1], **Chang Liu**[1,†], **Xintao Wang**[2], **Pengfei Wan**[2], **Kun Gai**[2], **Xiu Li**[1,†]
[1]Tsinghua University
[2]Kling Team, Kuaishou Technology
[3]Sun Yat-sen University

Project Page [§]Project Leader. [†]Corresponding Author. [*]Work Conducted During Internship.
liu-hl24@mails.tsinghua.edu.cn

Abstract

Recent advances in preference alignment for diffusion-based video generation, particularly via Direct Preference Optimization (DPO), have significantly improved visual quality. However, temporally sparse artifacts such as motion collapse, object flickering, and color oversaturation remain a major barrier to perceptual realism. Existing methods struggle with these issues due to two key limitations: (1) the preference attribution bottleneck, where offline human annotations are costly and fail to accurately capture learning dynamics, while online reward signals are rollout-aware but often unstable and biased; and (2) temporal credit misallocation, where uniformly applied supervision cannot effectively target the brief segments in which artifacts occur. To address these challenges, we propose concentrated Implicit Preference Optimization (cIPO), a post-training framework for video diffusion models. cIPO derives implicit preference signals directly from the denoising process: given a real video, the model adds forward noise and reconstructs it via iterative denoising, treating the original as the preferred sample and the reconstruction as the dispreferred one. This formulation captures inference-time errors without requiring human annotations or external reward models. Moreover, frame-level discrepancies between original and reconstructed videos reveal when failures occur. cIPO leverages this by computing temporal reconstruction errors and concentrating optimization on high-error segments, enabling more precise correction of failure-prone regions. Extensive experiments demonstrate that cIPO consistently enhances video authenticity and temporal coherence across multiple datasets, highlighting the effectiveness and efficiency of implicit preference with temporally concentrated optimization.

1 Introduction

Recent advances in diffusion models have substantially improved the visual quality[9, 8, 1, 2, 16, 17]. However, generating temporally coherent and artifact-free videos remains challenging. Unlike image generation, where perceptual failures are often spatially localized, video generation failures are often sparse in time: a video may appear plausible for most frames while exhibiting severe artifacts in only a few short segments, such as motion collapse, object flickering, or abrupt temporal discontinuities. These localized failures significantly degrade perceived video quality.

Preference optimization has recently emerged as a promising paradigm for post-training. By optimizing relative preferences between preferred and dispreferred outputs, methods such as direct preference optimization (DPO) offer an effective alternative to explicit reward modeling. However, applying

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

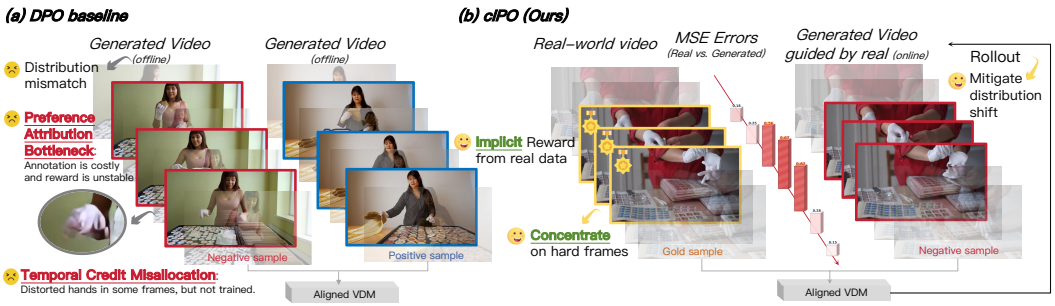

Figure 1: Compared with the baseline method, our method does not rely on additional human annotations or reward models. Besides, it can automatically identify difficult frames for training.

preference optimization to video generation presents two key challenges: preference attribution bottleneck and temporal credit misallocation.

The first challenge is the preference attribution bottleneck (*where reliable preference signals should come from*). As illustrated in Fig. 2, the core difficulty in video diffusion is not merely generating samples, but faithfully following the ideal reverse denoising trajectory back to the video data manifold [19, 14]. Starting from a perturbed real video (black arrow), an ideal denoising process should progressively recover the original sample and return to the data distribution (blue dashed arrow). In practice, however, due to limitations in model generative capacity, the learned denoising trajectory often deviates from this ideal path and accumulates errors over iterative rollout (red dashed arrow), leading to perceptual degradation in the final generated video. This rollout-induced deviation exposes a fundamental limitation of existing preference supervision. A straightforward solution is to collect human preference annotations by directly comparing generated videos. While such supervision is often high-quality, it is prohibitively expensive to scale and too static to track the model's evolving denoising behavior online. An alternative is to train reward models to score videos generated by the current policy. This captures rollout behavior but offers only indirect scalar signal on final outputs, and relies on external reward models, whose inherent biases may misalign with perceptual quality and induce optimization instability and reward hacking. The key issue is that preference signals should be both rollout-aware, so that they reflect the model's actual iterative generation behavior, and self-grounded, so that they arise from the model's own denoising dynamics rather than an external evaluator. This suggests a simple alternative paradigm: directly treat the discrepancy between the ideal target (real video) and the model's actual denoised result as an implicit preference signal, and optimize the model to reduce this rollout-induced drift (as shown by the purple arrows in Fig. 2).

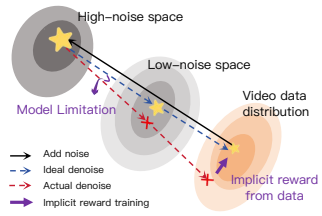

Figure 2: Ideal denoising should restore the video, but model trajectories drift due to error accumulation, naturally defining an implicit training preference.

The second challenge is temporal credit misallocation (*where optimization should be applied*). Even with reliable preference pairs, video preference optimization remains inefficient if optimization is distributed uniformly over time. We observe that errors in video diffusion are highly non-uniform (as shown in Fig. 3): A few frames exhibit significantly larger errors. Under such sparse failure patterns, uniform supervision wastes optimization on already well-generated frames while weakening learning on the temporally localized segments that matter most.

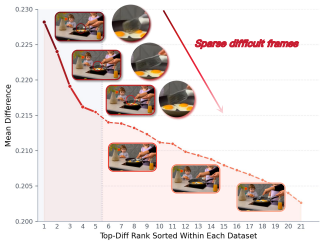

Figure 3: A few frames exhibit significantly larger errors.

To address these problems, we propose concentrated Implicit Preference Optimization (cIPO), a simple and effective framework for aligning video diffusion models. For preference attribution bottleneck, we constructs implicit preference supervision directly from diffusion denoising process. Given a real video, cIPO first perturbs it with forward diffusion noise, then reconstructs it through iterative denoising using the policy model online. This reconstruction process explicitly rolls out the model's own inference-time denoising trajectory, exposing the accumulated errors that emerge only during multi-step generation. The

discrepancy between the original video and its denoising result yields a naturally aligned preference signal: the original video serves as the preferred sample, while the reconstructed video serves as the dispreferred one. This formulation naturally aligns with the model's actual inference process, mitigating the training–inference mismatch without requiring human preference labels or external reward models.

To address temporal credit misallocation, cIPO further introduces concentrated preference optimization. Specifically, cIPO computes temporal reconstruction discrepancies, aggregates them over temporal windows, and allocates preference supervision only to the highest-error segments. By concentrating optimization on the temporal regions most responsible for perceptual degradation, cIPO substantially improves the density, precision, and effectiveness of preference optimizations.

Our work makes three contributions: (1) We identify preference attribution and temporal credit assignment as two central bottlenecks in video preference learning, demonstrating that video generation significantly exacerbates these challenges due to its iterative denoising trajectory and temporally sparse failure patterns. (2) We propose cIPO, a post-training framework that combines implicit denoised-based preference with concentrated temporal optimization. (3) We demonstrate that cIPO consistently improves video authenticity and temporal coherence across multiple benchmarks, showing that video preference learning benefits from implicit preference and concentrated temporal optimization.

2 Related Work

2.1 Preference Learning for Video Generation

Video diffusion preference optimization methods can be categorized by supervision source: human annotations, reward models, or synthetic data. Human-annotated methods directly optimize against pairwise labels. Flow-DPO[15] adapts DPO to video, while DenseDPO[23] uses dense segment-level preferences to reduce motion bias. These rely on costly offline annotations and are weakly aligned with current model failures. Reward-based methods use trained reward models. [15, 21, 21, 22, 27] score outputs for RL optimization. While these scale well and offer rollout-aware feedback, they are limited by training instability and reward hacking. Synthetic data methods avoid annotation by heuristically degrading real videos. DF-DPO[4] uses heuristic degradations (e.g., temporal reversal, frame shuffling) as negatives, LocalDPO[11] enhances this with localized corruptions, and RealDPO [3] directly contrasts real videos with model-generated outputs offline. However, it still does not explicitly capture how errors emerge during the denoising process itself. Overall, existing methods depend on costly or biased external reward, whereas cIPO leverages internal online denoising dynamics for precise, adaptive preferences.

2.2 Dynamic Temporal Sampling in Video Generation

Recent studies on dynamic sampling in video generation exploit temporal non-uniformity to achieve non-uniform computation allocation along the temporal dimension. A representative line of work includes DLFR-VAE [24], VGDFR [26], and DLFR-Gen [25], which observe that motion and information density vary substantially across time, and thus dynamically allocate fewer latent tokens to low-motion segments while preserving denser temporal representations for high-motion regions. Concretely, DLFR-VAE [24] introduces a dynamic latent frame-rate scheduler in the VAE latent space, while VGDFR [26] and DLFR-Gen [25] extend this idea to diffusion-based video generation via adaptive latent frame-rate scheduling and latent frame merging during denoising. Existing work dynamically redistributes computation for efficient generation, whereas our method dynamically redistributes alignment signal for temporally precise preference optimization, directly targeting sparse temporal artifacts such as flickering, motion collapse, and abrupt discontinuities.

3 Method

As shown in Fig. 4, cIPO framework consists of three components: (1) an implicit preference construction mechanism that derives preference pairs from reconstruction rollouts without external rewards; (2) a temporally concentrated selection strategy that focuses optimization on failure-prone

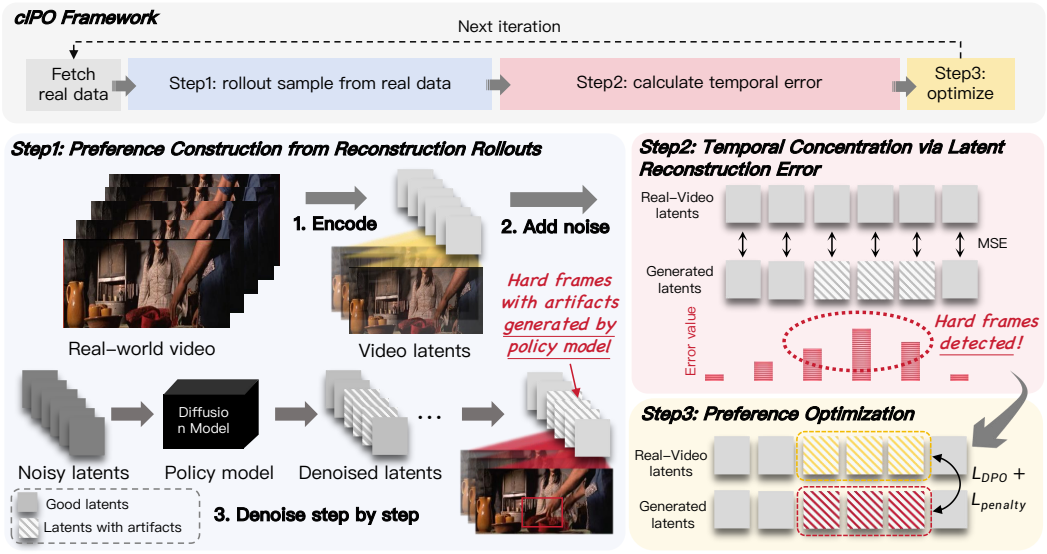

Figure 4: The figure shows the three-stage pipeline of cIPO. It first builds implicit preference pairs by reconstructing real videos through diffusion rollouts, using the original video as the preferred sample and its reconstruction as the dispreferred one. It then detects temporally localized hard frames via frame-wise latent reconstruction error, and finally concentrates DPO optimization on these high-error segments to improve temporal coherence and visual quality. Note that the entire process is carried out in the latent space. The videos in the figure are only for visual demonstration.

segments; and (3) a pairwise preference objective that improves temporal fidelity while preserving already-correct content.

3.1 Implicit Preference Construction from Reconstruction Rollouts

Let $x \in \mathcal{X}$ denote a real video sampled from the training distribution, and let $c \in \mathcal{C}$ denote its text condition. We encode x into the latent space of a pretrained video VAE encoder $\mathcal{E}$:

$$z_0 = \mathcal{E}(x) \in \mathbb{R}^{C\times T\times H\times W}, \tag{1}$$

where T is the number of latent frames. Let f_θ denote the trainable video diffusion model and $f_{\theta_{\text{ref}}}$ a frozen reference model. Given a diffusion starting index s, we construct a partially corrupted latent by applying the forward diffusion process to z_0. Specifically, let $\{t_k\}_{k=0}^{N-1}$ denote the scheduler timesteps and let $\epsilon \sim \mathcal{N}(0, I)$. We form $z_{t_s} = q(z_{t_s} \mid z_0, \epsilon)$, where q is the scheduler-induced forward noising process at timestep t_s. Starting from z_{t_s}, we run the reverse diffusion process conditioned on c using f_θ over the remaining timesteps. Denoting one reverse update by Φ_θ, we iteratively compute

$$z_{t_{k+1}} = \Phi_\theta(z_{t_k}, c, t_k), \qquad k = s, \dots, N-1, \tag{2}$$

and denote the terminal denoised latent by $\tilde{z}_0^{(s)}$.. We then define an implicit preference pair

$$y^+ = z_0, \qquad y^- = \tilde{z}_0^{(s)}, \tag{3}$$

where the clean latent is treated as the preferred sample and its noised-then-denoised reconstruction as the dispreferred sample. This yields one preference pair per training instance without requiring human annotations or a learned reward model. The specific algorithm process can refer to C.

We found that the reconstruction effect of the first frame is very poor. To reduce the first-frame drift and video structure collapse, we introduce a *clean first-frame anchoring* strategy. After obtaining y^-, we replace its first latent frame with that of the clean sample:

$$\hat{y}_\tau^- = \begin{cases} y_\tau^+, & \tau = 1, \\ y_\tau^-, & \tau > 1. \end{cases} \tag{4}$$

The anchored negative $\hat{y}^-$ preserves a stable reference for appearance and scene layout, while allowing later frames to expose temporal inconsistencies.

3.2 Temporal Concentration via Latent Reconstruction Error

A key observation is that temporal artifacts in generated videos are often concentrated in a few short segments. If preference optimization is applied uniformly over all frames, the gradients contributed by normal frames can overwhelm those from genuinely problematic segments, causing optimization to overfit temporally uninformative regions. cIPO addresses this issue by identifying the highest-error contiguous temporal window and restricting preference learning to that window.

For a preferred sample y^+ and an anchored negative $\hat{y}^-$, we define a framewise latent reconstruction discrepancy

$$d_\tau = \frac{1}{CHW} \left\| y_\tau^+ - \hat{y}_\tau^- \right\|_2^2, \qquad \tau = 1, \dots, T. \tag{5}$$

Here d_τ measures how severely the reconstruction deviates from the clean video at frame τ, T is the total number of frames in the original video. Given a window length K, we score each contiguous temporal window by its average discrepancy:

$$\mathcal{D}(a) = \frac{1}{K} \sum_{\tau=a}^{a+K-1} d_\tau, \qquad a \in \{1, \dots, T-K+1\}. \tag{6}$$

We then select the most failure-prone window

$$a^\star = \arg \max_{a \in \{1,\dots,T-K+1\}} \mathcal{D}(a), \qquad W^\star = \{a^\star, a^\star + 1, \dots, a^\star + K - 1\}. \tag{7}$$

Finally, we restrict both the preferred and dispreferred samples to the selected window:

$$\bar{y}^+ = y_{W^\star}^+, \qquad \bar{y}^- = \hat{y}_{W^\star}^-. \tag{8}$$

The use of a *contiguous* window is essential. Selecting scattered high-error frames independently may destroy short-range temporal structure and produce an incoherent optimization target that breaks short-range motion dependencies and weakens temporal consistency across adjacent frames. By contrast, Eq. (7) preserves local motion continuity and concentrates supervision on a semantically meaningful segment of the video.

3.3 Concentrated Implicit Preference Optimization

We now define the optimization objective. Let $\ell_\theta(y; c, t, \epsilon)$ denote the per-sample diffusion reconstruction loss of model f_θ on a latent video y, evaluated at diffusion step t with noise realization ϵ:

$$\ell_\theta(y; c, t, \epsilon) = \frac{1}{|\Omega|} \left\| f_\theta(y_t, c, t) - u(y, t, \epsilon) \right\|_2^2, \tag{9}$$

where $u(y, t, \epsilon)$ is the scheduler-specific regression target [15] and $|\Omega|$ is the number of latent elements. For a concentrated preferred sample $\bar{y}^+$ and a concentrated negative $\bar{y}^-$, we define the model-reference reconstruction gap

$$\Delta_\theta(y) = \ell_\theta(y; c, t, \epsilon) - \ell_{\theta_{\text{ref}}}(y; c, t, \epsilon). \tag{10}$$

Intuitively, $\Delta_\theta(y)$ measures whether the current model improves upon or degrades relative to the reference model on sample y. Given a concentrated preferred sample $\bar{y}^+$ and its corresponding concentrated negative $\bar{y}^-$, we define a pairwise preference logit

$$g = -\frac{\beta}{2} \left[\Delta_\theta(\bar{y}^+) - \Delta_\theta(\bar{y}^-) + \lambda \, \mathrm{ReLU}\big(\Delta_\theta(\bar{y}^+)\big) \right]. \tag{11}$$

where $\beta > 0$ is the preference sharpness coefficient and $\lambda \geq 0$ controls a stability regularizer.

The first term in Eq. (11) encourages the trainable model to achieve a larger relative improvement on the preferred sample than on the dispreferred one. The second term, $\mathrm{ReLU}(\Delta_\theta(\bar{y}^+))$, acts as a *winner-preservation penalty* following Smaug [18]: if the current model performs worse than the reference on the preferred clean segment, the penalty activates and suppresses this undesirable drift. This term is particularly important in the video setting, where aggressive preference updates can otherwise damage already-correct frames (see details in A).

The final cIPO objective adopts the standard DPO-style pairwise preference form $\mathcal{L}_{\text{cIPO}} = -\mathbb{E}_{(x,c)} \left[\log \sigma(g) \right]$, where $\sigma(\cdot)$ is the sigmoid function. This objective instantiates the standard DPO negative log-sigmoid loss with reconstruction-induced implicit preferences, encouraging the model to assign lower relative reconstruction error to the preferred clean segment than to its corresponding hard negative.

Table 1: Quantitative Comparison on MotionBench prompts from authenticity and motion dimensions. Best results are highlighted in **bold**, and second-best results are underlined. Forensic is the abbreviation of Forensic-chat, Om-D is the abbreviation of OmniAID-Dino, Off/On is the abbreviation of offline/online, and GT is the abbreviation of ground truth (i.e., real-world video).

Method	Frame Authenticity		Temporal Quality (Vbench)				
	Forensic ↑	Om-D ↑	Background ↑	Motion ↑	Subject ↑	Temporal ↑	Overall ↑
Pretrained	0.804	0.479	0.937	0.974	0.929	0.960	0.161
DPO (Off)	0.810	0.474	0.938	0.975	0.927	0.960	0.161
DPO (Off, GT)	0.815	0.478	<u>0.939</u>	0.972	0.919	0.957	0.162
DPO (On, GT)	0.830	<u>0.486</u>	0.931	0.973	0.919	0.958	<u>0.162</u>
DenseDPO	<u>0.819</u>	0.475	0.937	<u>0.975</u>	<u>0.930</u>	**0.963**	0.162
Ours	**0.876**	**0.524**	**0.947**	**0.989**	**0.937**	<u>0.961</u>	**0.168**

Table 2: Quantitative Comparison on WISA prompts from authenticity and motion dimensions. Best results are highlighted in **bold**, and second-best results are underlined.

Method	Frame Authenticity		Temporal Quality (VBench)				
	Forensic ↑	Om-D ↑	Background ↑	Motion ↑	Subject ↑	Temporal ↑	Overall ↑
Pre-trained	0.909	0.486	0.951	0.985	0.956	0.975	0.217
DPO (Off)	0.908	0.504	0.949	0.987	0.961	0.978	0.225
DPO (Off, GT)	0.916	0.491	**0.955**	0.985	0.960	0.976	0.222
DPO (On, GT)	<u>0.917</u>	<u>0.504</u>	0.951	0.987	0.961	0.979	<u>0.226</u>
DenseDPO	0.911	0.504	0.950	<u>0.987</u>	<u>0.962</u>	<u>0.979</u>	0.223
Ours	**0.931**	**0.521**	<u>0.952</u>	**0.990**	**0.989**	**0.983**	**0.247**

4 Experiment

4.1 Experimental Setup

Implementation details. We train on the training splits of MotionBench [6] and WISA [20], and evaluate on their held-out test prompts following [23]. MotionBench is a motion-centric benchmark for evaluating temporal reasoning and motion realism in text-to-video generation. WISA is a physics-aware text-to-video dataset covering diverse physical laws and scenarios. Since WISA does not provide an official split for preference learning, we construct our own train/test partition. Together, they enable evaluation of both motion fidelity and physics-aware temporal consistency. In order to compare the online construction methods of multiple types of negative samples, we instantiate our method on top of Wan-VACE [12], an all-in-one video foundation model that supports multiple conditional generation modes. To isolate the effect of preference optimization on motion generation, we restrict all experiments to the text-to-video (T2V) setting and optimize only the T2V generation pathway. We fine-tunes it using LoRA [7], targeting video generation at a resolution of 240×416, with learning rate of 5e-6, mixed-precision (bf16) training, gradient accumulation, and EMA. We set $\Delta = 0.1$ to perform regularization.

Compared methods. We compare our method against four representative baselines that differ in preference pair construction and supervision strategies: the pre-trained baseline refers to the original Wan-VACE checkpoint without preference optimization; DPO offline applies standard offline DPO using pairwise preferences derived from two model-generated videos under the same prompt, with preference labels determined by the VideoReward [15]; DPO offline (gt) introduces a variant that treats the ground-truth video as the positive sample and the generated video as the negative one, offering oracle supervision; DenseDPO[23] provides denser supervision by generating both positive and negative samples under real-video guidance and assigning preference scores at the temporal segment level. These baselines cover reward-model supervision, oracle preference construction, and dense temporal preference assignment, and therefore provide a representative comparison set for evaluation.

Table 3: Ablation on negative sample construction and sample strategy. We compare different negative sources and sampling strategies. For each negative source, selective sampling is compared against uniform sampling: ↑denotes improvement, ↓denotes degradation. Times(s) refers to the single-round sampling time. Best results are highlighted in **bold**, and second-best results are underlined.

Negative	Time (s)	Authenticity		VBench				
		Forensic ↑	Om-D ↑	Background ↑	Motion ↑	Subject ↑	Temporal ↑	Overall ↑
T2V	15	0.830	0.486	0.931	0.973	0.919	0.958	0.162
Frame	15	0.829	0.485	0.935	0.974	0.926	0.959	0.159
+ conc.	19	0.844↑	0.512↑	0.931↓	0.975↑	0.927↑	0.960↑	0.160↑
V2V	15	0.822	0.488	0.934	0.976	0.928	0.964	0.158
+ conc.	19	0.837↑	0.511↑	0.940↑	0.978↑	0.931↑	0.966↑	0.157↓
Noise	4	0.862	0.494	0.938	0.979	0.925	0.966	0.157
+ conc.	6	**0.876**↑	**0.524**↑	**0.947**↑	**0.989**↑	**0.937**↑	**0.961**↓	**0.168**↑

Evaluation metrics. We evaluate generated videos from two complementary perspectives: authenticity and motion quality. For authenticity, we report Forensic-Chat [13], a forensic realism score assessing visual authenticity and synthetic artifact suppression, along with OmniAID-Dino [5], an authenticity metric based on DINO that measures semantic realism and naturalness. Regarding motion quality, we follow the VBench [10] protocol and report scores for background consistency, motion smoothness, subject consistency, temporal flickering and overall consistency.

4.2 Quantitative Results

Table 1 and Table 2 present the quantitative results on the MotionBench and Wisa test sets, respectively. Our method consistently outperforms all baselines across both authenticity and motion-related metrics, demonstrating its effectiveness in improving perceptual realism while preserving motion fidelity. Compared with DenseDPO, our method yields a substantial gain in authenticity ($0.819/0.475 \rightarrow 0.876/0.524$), indicating that directly leveraging real videos as supervision provides a more reliable training signal than reward-model-based dense supervision that scores individual temporal segments.

To disentangle the contribution of real videos as positive samples from the improvement introduced by our analysis framework itself, we conduct DPO (Online,GT) in which the ground-truth video is used as the positive sample and the generated video as the negative sample for both methods . Under this setting, our method still improves Forensic from 0.830 to 0.876 and OmniAID-D from 0.486 to 0.524. These gains indicate that the performance improvement is not solely due to the introduction of real data, but also arises from our preference attribution strategy. In particular, these results suggest that deriving supervision from the model's own denoising trajectory and temporal concentration is an effective strategy.

4.3 Qualitative Results

As shown in Fig. 5, compared with global preference optimization methods such as DPO and DenseDPO, cIPO produces markedly more temporally coherent videos across both fine-grained manipulation and authenticity. Specifically, in challenging motion regions highlighted by the red boxes, DPO and DenseDPO often exhibit noticeable temporal inconsistency, motion blur, and structural distortions, such as unstable hand-object interactions, inconsistent body poses during spinning motions, and unrealistic human motion trajectories. By contrast, cIPO maintains smoother temporal transitions, more stable object geometry, and more realistic motion dynamics throughout the video sequence. In contrast, cIPO leverages denoised-base implicit preference signals to identify failure-prone moments and concentrates optimization on high-error temporal windows, resulting in substantially improved authenticity and temporal stability.



Figure 5: Qualitative Results. Compared to baselines, cIPO achieves better temporal coherence in challenging motion regions highlighted by the red boxes.

4.4 Ablation Studies

We ablate the two core design choices in our method: how negative samples are constructed and how preference supervision is temporally allocated. Table 3 summarizes the results, while Fig. 6 further analyze noise starting step impact in our framework.

Effect of negative sample from reconstruction rollouts. For negative sample construction, we leverage VACE to construct structurally aligned negatives through two controlled generation strategies: (1) first-frame-conditioned generation, in which negative samples share the identical initial latent as positive samples, and (2) video-to-video (V2V) generation, where the negative is generated under real-video guidance to remain close to the positive trajectory. For a fair comparison, all methods are tested under the same GPU-hour training budget. Similar to prior findings, these aligned pair construction strategies yield moderate improvements over unconstrained T2V negatives, particularly on motion-related metrics (as shown in Table. 3). However, reconstruction-based noise negatives consistently outperform all alternative negative sources. This result supports our central design: negatives induced by diffusion perturbation are more informative than externally constructed samples because they better expose the model's own inference-time failure modes. In addition, since denoise only requires a small number of sampling steps (about 5 to 20 steps), while other methods require complete sampling (50 steps), the training time and cost of the denoise method are lower.

Effect of temporal concentration strategy. To validate the temporal concentration strategy, we apply it to all negative sample variants (excluding T2V due to substantial disparity between positive and negative samples that hinder direct comparison). As shown in Table. 3, the consistent performance gains confirm its orthogonality with respect to the negative sampling source. This demonstrates the effectiveness and universality of the temporal concentration strategy.

Noise starting step. Fig. 6 examines the diffusion starting index sfor constructing reconstruction negatives. Results show that moderate noise levels work best: too little noise yields overly simple negatives, while excessive noise disrupts reconstruction and semantic consistency. This observation is consistent with the intuition that useful preference pairs should be difficult enough to expose rollout failure, but not so corrupted as to become semantically uninformative.

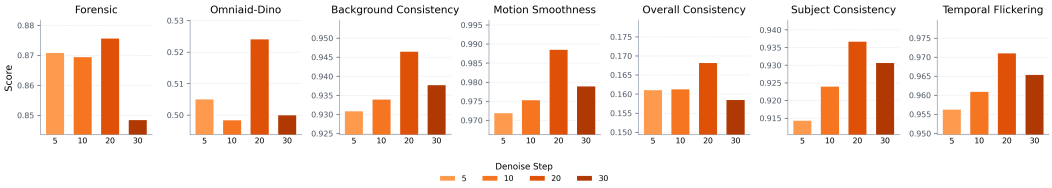

Figure 6: The impact of the change of the starting index s of noise addition used in the negative sample sampling process on the video temporal quality and authenticity.

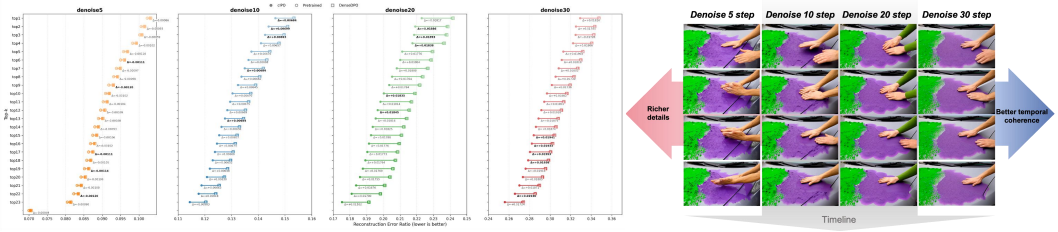

(a) Analysis of Reconstruction Errors in the Inference Process

(b) Visualization of the generation effects

Figure 7: Analysis of the effect of different noise levels

4.5 Analysis

Reconstruction Errors in the Inference Process. Fig. 7.a provides a comprehensive validation of the two core design choices in cIPO: implicit preference construction and temporally concentrated optimization. The x-axis shows the Reconstruction Error Ratio (lower is better), while the y-axis denotes the top-k highest-error temporal windows selected for preference optimization. The four subplots correspond to different forward noise levels, where larger denoise steps indicate stronger forward perturbation and thus a more challenging reconstruction task. Across all noise levels, cIPO consistently outperforms Pretrained and DenseDPO, demonstrating that constructing preference pairs from the original video and its denoised reconstruction provides a more effective supervision signal. Specifically, under a moderate number of noise level (e.g., 10 and 20), as shown in Fig. 6, the model achieves the best authenticity and motion quality trade-off. This is likely because most severely corrupted frames are corrected (as illustrated in Fig. 7.a, the most significant improvements are concentrated in the first four steps for the 10-noise and 20-noise settings), leading to substantial performance gains, which further validates the effectiveness of our temporal concentration strategy.

Visualization of the generation effects of models. The generation results under different noise levels reflect the natural trade-off between authenticity and temporal consistency. As shown in Fig. 7.b, under weak noise, the reconstruction results can better preserve local textures and spatial details, but they are insufficient in exposing temporal artifacts, and local jitter and inter-frame flicker are still likely to occur. Under strong noise, although the model can learn smoother motion patterns and more stable cross-frame consistency, it often comes with detail blurring and texture loss. This phenomenon indicates that cIPO lies in constructing preference pairs using noise of appropriate level, enabling the model to achieve a better compromise between visual details and temporal consistency.

5 Conclusions and Limitations

In this work, we propose concentrated Implicit Preference Optimization (cIPO), a post-training framework for improving temporal coherence in text-to-video generation. By deriving preference signals directly from the model's denoising process, cIPO avoids the need for costly human annotations or unstable external reward models, enabling efficient and rollout-consistent preference learning. Moreover, cIPO identifies temporally localized failure regions and concentrates optimization on high-error segments, leading to more precise training for short but perceptually critical artifacts. Experimental results show that cIPO consistently improves both authenticity and temporal consistency across multiple datasets, demonstrating the effectiveness of denoised-based preference with temporally focused optimization. However, cIPO still relies on reconstruction error as a proxy for perceptual quality, which may not fully capture high-level semantic preferences or human subjective judgments.

Acknowledgments and Disclosure of Funding

This work was supported in part by the National Natural Science Foundation of China (Grants 62406167 and U24B6012), the Shenzhen Key Laboratory of New Generation Interactive Media Technology Innovation (Grant No. ZDSYS20210623092001004) and the Kling Team at Kuaishou Technology.

References

[1] Chubin Chen, Sujie Hu, Jiashu Zhu, Meiqi Wu, Jintao Chen, Yanxun Li, Nisha Huang, Chengyu Fang, Jiahong Wu, Xiangxiang Chu, et al. Taming preference mode collapse via directional decoupling alignment in diffusion reinforcement learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12775–12786, 2026.

[2] Chubin Chen, Jiashu Zhu, Xiaokun Feng, Nisha Huang, Chen Zhu, Meiqi Wu, Fangyuan Mao, Jiahong Wu, Xiangxiang Chu, and Xiu Li. Stochastic self-guidance for training-free enhancement of diffusion models. In *International Conference on Learning Representations*, volume 2026, pages 134462–134490, 2026.

[3] Guo Cheng, Danni Yang, Ziqi Huang, Jianlou Si, Chenyang Si, and Ziwei Liu. Realdpo: Real or not real, that is the preference. *arXiv preprint arXiv:2510.14955*, 2025.

[4] Haoran Cheng, Qide Dong, Liang Peng, Zhizhou Sha, Weiguo Feng, Jinghui Xie, Zhao Song, Shilei Wen, Xiaofei He, and Boxi Wu. Discriminator-free direct preference optimization for video diffusion. *arXiv preprint arXiv:2504.08542*, 2025.

[5] Yuncheng Guo, Junyan Ye, Chenjue Zhang, Hengrui Kang, Haohuan Fu, Conghui He, and Weijia Li. Omniaid: Decoupling semantic and artifacts for universal ai-generated image detection in the wild. *arXiv preprint arXiv:2511.08423*, 2025.

[6] Wenyi Hong, Yean Cheng, Zhuoyi Yang, Weihan Wang, Lefan Wang, Xiaotao Gu, Shiyu Huang, Yuxiao Dong, and Jie Tang. Motionbench: Benchmarking and improving fine-grained video motion understanding for vision language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 8450–8460, 2025.

[7] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Liang Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *Iclr*, 1(2):3, 2022.

[8] Sujie Hu, Chubin Chen, Jiashu Zhu, Jiahong Wu, Xiangxiang Chu, and Xiu Li. Embedding-perturbed exploration preference optimization for flow models. In *Forty-third International Conference on Machine Learning*, 2026.

[9] Sujie Hu, Huaisong Zhang, Pujing Yang, Qian Wang, Xiaomin Li, Xinhang Leng, Chen Li, Jing Lyu, Xiu Li, and Ling Pan. Blame errant words, not the canvas: Defect-aware spatial credit assignment for visual text generation, 2026.

[10] Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, et al. Vbench: Comprehensive benchmark suite for video generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21807–21818, 2024.

[11] Zitong Huang, Kaidong Zhang, Yukang Ding, Chao Gao, Rui Ding, Ying Chen, and Wangmeng Zuo. Mind the generative details: Direct localized detail preference optimization for video diffusion models. *arXiv preprint arXiv:2601.04068*, 2026.

[12] Zeyinzi Jiang, Zhen Han, Chaojie Mao, Jingfeng Zhang, Yulin Pan, and Yu Liu. Vace: All-in-one video creation and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17191–17202, 2025.

[13] Kaiqing Lin, Zhiyuan Yan, Ruoxin Chen, Junyan Ye, Ke-Yue Zhang, Yue Zhou, Peng Jin, Bin Li, Taiping Yao, and Shouhong Ding. Seeing before reasoning: A unified framework for generalizable and explainable fake image detection. *arXiv preprint arXiv:2509.25502*, 2025.

[14] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2022.

[15] Jie Liu, Gongye Liu, Jiajun Liang, Ziyang Yuan, Xiaokun Liu, Mingwu Zheng, Xiele Wu, Qiulin Wang, Menghan Xia, Xintao Wang, et al. Improving video generation with human feedback. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[16] Zijie Meng. Orthomotion: Disentangling camera and subject motion via geometry semantics orthogonal attention. *arXiv preprint arXiv:2606.22835*, 2026.

[17] Zijie Meng, Jiwen Liu, Yufei Liu, Chengzhuo Tong, Xiaoqiang Liu, Yuanxing Zhang, Yulong Xu, and Pengfei Wan. Argus: Stacked multi-view identity mosaic injection for subject-preserving video generation. *arXiv preprint arXiv:2606.11670*, 2026.

[18] Arka Pal, Deep Karkhanis, Samuel Dooley, Manley Roberts, Siddartha Naidu, and Colin White. Smaug: Fixing failure modes of preference optimisation with dpo-positive. *arXiv preprint arXiv:2402.13228*, 2024.

[19] 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*, 2020.

[20] Jing Wang, Ao Ma, Ke Cao, Jun Zheng, Jiasong Feng, Zhanjie Zhang, Wanyuan Pang, and Xiaodan Liang. Wisa: World simulator assistant for physics-aware text-to-video generation. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[21] Yibin Wang, Zhiyu Tan, Junyan Wang, Xiaomeng Yang, Cheng Jin, and Hao Li. Lift: Leveraging human feedback for text-to-video model alignment. *arXiv preprint arXiv:2412.04814*, 2024.

[22] Yibin Wang, Yuhang Zang, Hao Li, Cheng Jin, and Jiaqi Wang. Unified reward model for multimodal understanding and generation. *arXiv preprint arXiv:2503.05236*, 2025.

[23] Ziyi Wu, Anil Kag, Ivan Skorokhodov, Willi Menapace, Ashkan Mirzaei, Igor Gilitschenski, Sergey Tulyakov, and Aliaksandr Siarohin. Densedpo: Fine-grained temporal preference optimization for video diffusion models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[24] Zhihang Yuan, Siyuan Wang, Yuzhang Shang, Hanling Zhang, Tongcheng Fang, Rui Xie, Shengen Yan, Guohao Dai, and Yu Wang. Dlfr-vae: Dynamic latent frame rate vae for video generation. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 10388–10397, 2025.

[25] Zhihang Yuan, Rui Xie, Yuzhang Shang, Hanling Zhang, Siyuan Wang, Shengen Yan, Guohao Dai, and Yu Wang. Dlfr-gen: Diffusion-based video generation with dynamic latent frame rate. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16410–16419, October 2025.

[26] Zhihang Yuan, Rui Xie, Yuzhang Shang, Hanling Zhang, Siyuan Wang, Shengen Yan, Guohao Dai, and Yu Wang. Vgdfr: Diffusion-based video generation with dynamic latent frame rate. *arXiv preprint arXiv:2504.12259*, 2025.

[27] Jiacheng Zhang, Jie Wu, Weifeng Chen, Yatai Ji, Xuefeng Xiao, Weilin Huang, and Kai Han. Onlinevpo: Align video diffusion model with online video-centric preference optimization. *arXiv preprint arXiv:2412.15159*, 2024.

A Analysis of Gradient Degeneration with Near-Identical Pairs

The design of winner-preservation penalty is primarily motivated by a key failure mode in Flow-DPO: when the preferred and dispreferred trajectories are similar, the probability of the preferred sample can actually decrease. In this section, we analyze the causes of this problem.

Setup Consider a flow model parameterized by θ with velocity predictor $v_\theta(x_t, t)$. For a trajectory sample (x_t, v) at time t, the conditional likelihood is

$$\log p_\theta(v \mid x_t, t) = -\frac{\beta_t}{2}\|v - v_\theta(x_t, t)\|^2 + C,$$

where $\beta_t > 0$ is a time-dependent weighting coefficient and C is independent of θ.

Given a preferred trajectory w and a dispreferred trajectory l, Flow-DPO optimizes the objective

$$\mathcal{L}_{\text{FDPO}} = -\mathbb{E}\left[\log \sigma\left(\Delta_\theta\right)\right],$$

where

$$\Delta_\theta = \log p_\theta(v^w \mid x_t^w, t) - \log p_\theta(v^l \mid x_t^l, t) - \left(\log p_{\text{ref}}(v^w \mid x_t^w, t) - \log p_{\text{ref}}(v^l \mid x_t^l, t)\right).$$

Since the reference model is fixed, the optimization dynamics are governed by

$$\Delta_\theta = -\frac{\beta_t}{2}\left(\|v^w - v_\theta(x_t^w, t)\|^2 - \|v^l - v_\theta(x_t^l, t)\|^2\right) + \text{const}.$$

We now formalize the causal relationship between gradient competition and the decrease of preferred likelihood in Flow-DPO.

Recall that the Flow-DPO update direction is

$$\delta\theta \propto g_w - g_l,$$

where

$$g_w = \nabla_\theta \log p_\theta(v^w \mid x_t^w, t), \qquad g_l = \nabla_\theta \log p_\theta(v^l \mid x_t^l, t).$$

To analyze how this update affects the preferred trajectory, we consider the first-order change of the preferred log-likelihood:

$$\delta \log p_\theta(v^w \mid x_t^w, t) \approx g_w^\top \delta\theta.$$

Substituting the Flow-DPO update gives

$$\delta \log p_\theta(v^w \mid x_t^w, t) \propto g_w^\top (g_w - g_l).$$

Expanding the inner product yields

$$\delta \log p_\theta(v^w \mid x_t^w, t) \propto \underbrace{\|g_w\|^2}_{\text{preferred enhancement}} - \underbrace{g_w^\top g_l}_{\text{dispreferred suppression}}.$$

This decomposition reveals two competing forces:

- The term $\|g_w\|^2$ corresponds to the standard likelihood-increasing effect for the preferred trajectory.
- The term $g_w^\top g_l$ arises from suppressing the dispreferred trajectory and measures how much this suppression interferes with the preferred update direction.

Similarity induces gradient alignment. When the preferred and dispreferred trajectories are temporally similar, i.e.,

$$x_t^w \approx x_t^l, \qquad t_w \approx t_l,$$

their network Jacobians become highly correlated:

$$\nabla_\theta v_\theta(x_t^w, t) \approx \nabla_\theta v_\theta(x_t^l, t).$$

Consequently, the corresponding likelihood gradients become aligned:

$$g_w^\top g_l \approx \|g_w\|\,\|g_l\|.$$

Why the dispreferred suppression term may dominate. After Flow Matching pretraining or supervised finetuning, the preferred trajectory is typically already well fitted:

$$\|v^w - v_\theta(x_t^w, t)\| \ll \|v^l - v_\theta(x_t^l, t)\|.$$

Using

$$g = \nabla_\theta \log p_\theta(v \mid x_t, t) = \beta_t (v - v_\theta(x_t, t))^\top \nabla_\theta v_\theta(x_t, t),$$

the gradient magnitude approximately scales with the residual norm:

$$\|g\| \propto \|v - v_\theta(x_t, t)\|.$$

Therefore,

$$\|g_l\| > \|g_w\|.$$

Combined with gradient alignment, this implies

$$g_w^\top g_l > \|g_w\|^2.$$

Hence,

$$\delta \log p_\theta(v^w \mid x_t^w, t) < 0.$$

Equivalently,

$$p_\theta(v^w \mid x_t^w, t)$$

decreases after the Flow-DPO update.

Interpretation. The key issue is that Flow-DPO optimizes a *relative preference objective* rather than directly maximizing the absolute likelihood of preferred trajectories. Under strong overlap, suppressing the dispreferred trajectory requires parameter updates that are highly aligned with the preferred trajectory gradients. When the dispreferred gradient magnitude is larger, the suppression effect dominates the enhancement effect, causing the preferred likelihood to decrease despite optimizing a preference objective. In practice, this often manifests as spurious high-frequency textures or over-smoothed motion patterns. To address this limitation, we add a *winner-preservation penalty*, which is activated when the current model underperforms the reference model on the preferred clean segment.

Experiment. In flow matching, the model learns a conditional velocity field $v_\theta(x_t, c, t)$ by minimizing the regression objective

$$\mathcal{L}_{\mathrm{FM}} = \mathbb{E}\big[\|v_\theta(x_t, c, t) - u_t\|^2\big],$$

where u_t denotes the target velocity. Under the standard Gaussian regression interpretation,

$$p_\theta(u_t \mid x_t, c, t) \propto \exp\left(-\frac{1}{2}\|v_\theta - u_t\|^2\right),$$

which implies

$$\log p_\theta(u_t \mid x_t, c, t) = -\frac{1}{2}\|v_\theta - u_t\|^2 + \text{const}.$$



Figure 8: Variation of w_{diff} during the training process with/without the penalty term.

Table 4: Percentages are computed within each dimension

Criterion	Pretrained	DenseDPO	DPO	cIPO
Human Structure	11.1%	11.1%	14.8%	63.0%
Motion Dynamics	6.9%	6.9%	20.7%	65.5%
Scene Structure	13.8%	10.3%	17.2%	58.6%
Object Relations	10.3%	6.9%	13.8%	69.0%
Natural Physics	10.3%	3.4%	13.8%	72.4%
Temporal Consistency	10.3%	10.3%	10.3%	69.0%

Therefore, for a preferred sample x_w,

$$w_{\text{diff}} = \texttt{model_win_err} - \texttt{ref_win_err}$$

corresponds to the relative log-likelihood change between the current model and the reference model:

$$\log \frac{p_\theta(x_w)}{p_{\text{ref}}(x_w)} \propto -\frac{1}{2}\, w_{\text{diff}}.$$

Consequently, a larger positive w_{diff} indicates that the current model assigns a lower probability to the preferred sample compared to the reference model.

As shown in Fig. 8, without the penalty term, w_{diff} increases significantly during the later stage of training, indicating that the model progressively reduces the likelihood of preferred samples. This reveals a preference degeneration phenomenon where optimization unintentionally suppresses positive samples themselves. In contrast, with the proposed penalty, w_{diff} remains close to zero or negative, demonstrating that the preferred-sample probability is preserved or improved throughout training. These results show that the penalty stabilizes preference optimization by preventing the collapse of preferred-sample likelihood.

B Human evaluation

We conducted a human evaluation to compare four models: *Pretrained*, *DenseDPO*, *DPO*, and *cIPO*. For each sample, annotators were asked to select the best model output along six criteria: *human structure*, *motion dynamics*, *scene structure*, *object relations*, *natural physics*, and *temporal consistency*. The spreadsheet records the majority-vote winner for each sample–criterion pair. To obtain a dimension-level ranking, we counted how many samples were won by each model under the majority vote for that criterion. The resulting comparison is therefore based on majority-vote wins across samples rather than mean opinion scores. As summarized in Table 4, *cIPO* is the strongest model on all six criteria, consistently achieving the highest majority-vote share for every dimension. The advantage of *cIPO* is especially pronounced on *natural physics* (72.4%) and *object relations / temporal consistency* (69.0% each), while the other three models remain substantially behind on every dimension.

C Pseudocode for Rollout

The pseudocode 1 shows the specific construction method of negative samples in Implicit Preference Construction from Reconstruction Rollouts.

Algorithm 1 Rollout of a Noisy Latent Trajectory

Require: $\mathbf{z}_0$, scheduler $\mathcal{S}$, denoiser $\mathcal{M}$, prompt embedding $\mathbf{c}$, negative prompt $\mathbf{c}^-$, steps N, guidance w, start step s
Ensure: Denoised latent $\hat{\mathbf{z}}_0$
$\{t_k\}_{k=0}^{N-1} \leftarrow \mathcal{S}.\texttt{set_timesteps}(N)$
$\{t_k\}_{k=s}^{N-1} \leftarrow$ active steps
$\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
$\mathbf{z}_{t_s} \leftarrow \mathcal{S}.\texttt{add_noise}(\mathbf{z}_0, \boldsymbol{\epsilon}, t_s)$
for $t_k \in \{t_k\}_{k=s}^{N-1}$ **do**
  $\hat{\boldsymbol{\epsilon}}_{\text{cond}} \leftarrow \mathcal{M}(\mathbf{z}_{t_k}, t_k, \mathbf{c})$
  if $w > 1$ and $\mathbf{c}^-$ provided **then**
    $\hat{\boldsymbol{\epsilon}}_{\text{uncond}} \leftarrow \mathcal{M}(\mathbf{z}_{t_k}, t_k, \mathbf{c}^-)$
    $\hat{\boldsymbol{\epsilon}} \leftarrow \hat{\boldsymbol{\epsilon}}_{\text{uncond}} + w\big(\hat{\boldsymbol{\epsilon}}_{\text{cond}} - \hat{\boldsymbol{\epsilon}}_{\text{uncond}}\big)$
  else
    $\hat{\boldsymbol{\epsilon}} \leftarrow \hat{\boldsymbol{\epsilon}}_{\text{cond}}$
  end if
  $\mathbf{z}_{t_{k+1}} \leftarrow \mathcal{S}.\texttt{step}(\hat{\boldsymbol{\epsilon}}, t_k, \mathbf{z}_{t_k})$
end for
$\hat{\mathbf{z}}_0 \leftarrow \mathbf{z}_{t_N}$
return $\hat{\mathbf{z}}_0$

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [N/A].
- [N/A] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for [N/A]).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will also be asked to include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While [Yes] is generally preferable to [No], it is perfectly acceptable to answer [No] provided a proper justification is given (e.g., error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering [No] or [N/A] is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist"**,
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers**.

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: In Method section.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: In limitations section.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: In Appendix section.

   Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: In Experiement section.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: In Experiement section.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: In Experiement section.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: The training cost is relatively high.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.

- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: In Experiement section.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: In Experiement section.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: In Conclusion section.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: In Conclusion section.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: In Experiment section.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: No new assets.

Guidelines:

- The answer [N/A] means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [Yes]

    Justification: In Appendix section.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: In Conclusion section.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLM only for writing, editing, or formatting purposes.

    Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.