

# Ouroboros-Diffusion: Exploring Consistent Content Generation in Tuning-free Long Video Diffusion\*

Jingyuan Chen<sup>1</sup>, Fuchen Long<sup>2</sup>, Jie An<sup>1</sup>, Zhaofan Qiu<sup>2</sup>, Ting Yao<sup>2</sup>, Jiebo Luo<sup>1</sup>, Tao Mei<sup>2</sup>

<sup>1</sup> University of Rochester, Rochester, NY USA

<sup>2</sup> HiDream.ai Inc.

jchen157@u.rochester.edu, longfuchen@hidream.ai, jan6@cs.rochester.edu  
{qiuzhaofan, tiyao}@hidream.ai, jluo@cs.rochester.edu, tmei@hidream.ai

## Abstract

The first-in-first-out (FIFO) video diffusion, built on a pre-trained text-to-video model, has recently emerged as an effective approach for tuning-free long video generation. This technique maintains a queue of video frames with progressively increasing noise, continuously producing clean frames at the queue’s head while Gaussian noise is enqueued at the tail. However, FIFO-Diffusion often struggles to keep long-range temporal consistency in the generated videos due to the lack of correspondence modeling across frames. In this paper, we propose Ouroboros-Diffusion, a novel video denoising framework designed to enhance structural and content (subject) consistency, enabling the generation of consistent videos of arbitrary length. Specifically, we introduce a new latent sampling technique at the queue tail to improve structural consistency, ensuring perceptually smooth transitions among frames. To enhance subject consistency, we devise a Subject-Aware Cross-Frame Attention (SACFA) mechanism, which aligns subjects across frames within short segments to achieve better visual coherence. Furthermore, we introduce self-recurrent guidance. This technique leverages information from all previous cleaner frames at the front of the queue to guide the denoising of noisier frames at the end, fostering rich and contextual global information interaction. Extensive experiments of long video generation on the VBench benchmark demonstrate the superiority of our Ouroboros-Diffusion, particularly in terms of subject consistency, motion smoothness, and temporal consistency.

## 1 Introduction

With the rise of artificial visual content generation technologies, significant breakthroughs have been made in video diffusion (Blattmann et al. 2023; Peng et al. 2024; Ma et al. 2024b; Long et al. 2024). However, most current video diffusion models are trained on short clips (*e.g.*, 16 frames), a significant challenge when scaling to long video generation. Instead of relying on extensive training data to extend video length, our work focuses on leveraging a pre-trained video diffusion model to generate long videos without extensive training data or fine-tuning. Additionally, we place a strong emphasis on content consistency to enhance both the visual and motion quality of long videos in diffusion.

\*This work was performed at HiDream.ai.  
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

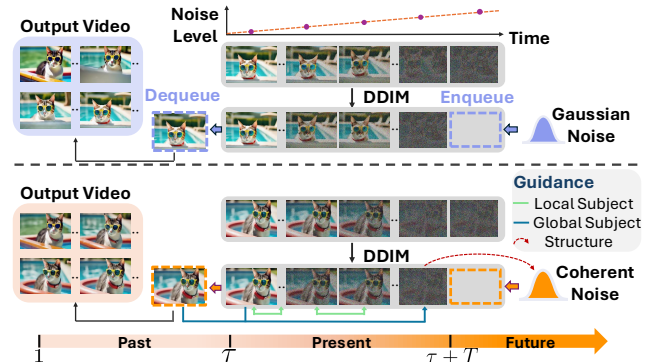


Figure 1: Illustration of FIFO-Diffusion (Kim et al. 2024) (top) and our Ouroboros-Diffusion (bottom) for tuning-free long video generation.

The recent first-in-first-out diffusion (Kim et al. 2024) has successfully employed the pre-trained video diffusion model for infinite frame generation. Compared to general video diffusion that all frames share the same noise level, the diagonal denoising approach proposed by Kim et al. (2024) maintains a queue of frames with progressively increasing noise levels, enabling frame-by-frame generation at each step. The upper part of Figure 1 conceptualizes the diagonal denoising process. In this process, a fully denoised frame at the queue head is popped out, while a random noisy latent is pushed to the tail. The latent enqueue-dequeue cycle allows for the incremental generation of video frames. However, the independently enqueued Gaussian noise can lead to content discrepancies between the frame latents near the tail of the queue. Moreover, FIFO-Diffusion fails to utilize visual information contained in the earlier denoised frames, which exacerbates temporal inconsistency in continuous video diffusion. For instance, the appearance of the cat changes significantly in the output video of Figure 1. Our work addresses this issue with FIFO-Diffusion from two key perspectives: structural and subject consistency. To this end, we introduce a novel denoising framework termed Ouroboros-Diffusion, designed for tuning-free long video generation. Inspired by the ancient symbol of Ouroboros—a serpent or dragon eating its own tail, symbolizing wholeness and self-renewal—our framework embodies these concepts by seam-

lessly integrating information across time. The design of Ouroboros-Diffusion is guided by three core principles targeting distinct information flows to improve structural and subject consistency: *present infers future*, *present influences present*, and *past informs present*, as depicted in Figure 1. To enhance structural consistency, we address the critical step of enqueueing new tail noise, which can lead to structural incoherence. From a video continuation perspective, this step involves sampling a *future* frame that should maintain visual structure continuity with previous frames. To achieve this, we propose inferring the *future* frame from the *present* frame in the denoising queue by exploiting the low-frequency component connection between them. Specifically, instead of initializing the tail latent with Gaussian noise, Ouroboros-Diffusion extracts the low-frequency component from the second-to-last frame latent using Fast Fourier Transform (FFT) and combines it with the high-frequency part of random noise to create the enqueued latent. The low-frequency component preserves layout information for overall video consistency, while the high-frequency component introduces necessary video dynamics. For subject consistency, we consider the semantic dependencies of long videos from two angles: within the current denoising queue and between previously generated clear frames and the current queue. To enhance subject temporal coherence within the *present* queue, we extend self-attention across frames through a Subject-Aware Cross-Frame Attention (SACFA) module. This module leverages segmented subject regions from the cross-attention map to extract subject tokens in each frame, which serve as auxiliary contexts for subject alignment. These tokens are then stored in a subject feature bank. To model longer-range subject dependencies, we introduce a long-term memory of *past* subjects to guide the appearance of the *present* subject. Specifically, Ouroboros-Diffusion utilizes the long-term memory derived from the frame at the head of the queue to guide the denoising of noisier frames near the tail, optimizing the latent through a subject-aware gradient during video denoising.

The main contribution of this work is the proposal of Ouroboros-Diffusion to address content consistency in turning-free long video generation. Our solution elegantly explored how diagonal denoising could benefit from low-frequency content preservation, and how to preserve subject consistency through cross-frame attention and gradient-based latent optimization in diffusion. Extensive experiments on VBench verify the effectiveness of our proposal in terms of both visual and motion quality.

## 2 Related Work

**Text-to-Video Diffusion Models.** The great success of text-to-video (T2V) diffusion models (Ho et al. 2022a; Voleti, Jolicoeur-Martineau, and Pal 2022; Villegas et al. 2023; Wang et al. 2023b; Yin et al. 2023a; Chen et al. 2023; Zhang et al. 2024; Guo et al. 2024) for video generation based on text prompts has been witnessed in recent years. VDM (Ho et al. 2022b) is one of the early works that combines spatial and temporal attention to construct space-time factorized UNet for video synthesis. Later in Make-A-Video (Singer et al. 2023), the prior knowledge of text-to-image diffusion

models is explored in video diffusion and the 2D-UNet is extended with the temporal modules (*e.g.*, temporal convolution and self-attention (Long et al. 2022a)) for motion modeling (Long et al. 2019, 2022b, 2023). The advances (An et al. 2023; Blattmann et al. 2023; Wu et al. 2023a; Chen et al. 2024b) further execute the video synthesis on latent space and push the boundaries of high-resolution video generation. Inspired by the impressive performances of Diffusion Transformers (DiTs) (Peebles and Xie 2023; Ma et al. 2024a), the spatial-temporal transformer architecture starts to emerge in video diffusion (Hong et al. 2023; Xu et al. 2024). Here, we choose VideoCrafter2 (Chen et al. 2024a) as a backbone text-to-video model for long video generation.

**Long Video Diffusion.** Despite the achievements of text-to-video diffusion, long video generation is still a grand challenge. Existing works have explored two strategies, *i.e.*, tuning-based and tuning-free long video diffusion. Typically, the tuning-based methods (Yin et al. 2023b; Henschel et al. 2024; Tian et al. 2024; Jin et al. 2024) usually exploits an auto-regressive manner which leverages the information of past generated frames to guide the synthesis of current frames. Nevertheless, tuning the auto-regressive video diffusion model usually involves a huge computational cost. To overcome this limitation, tuning-free approaches (Wang et al. 2023a; Qiu et al. 2024; Kim et al. 2024; Oh et al. 2024; Tan et al. 2024) try to adopt a pre-trained basic video diffusion model with the power of short-clip synthesis to generate multiple frames with temporal coherence. Qiu et al. (2024) introduces a sliding-window temporal attention mechanism to simultaneously denoising all frames for keeping temporal consistency. Recent advance FIFO-Diffusion (Kim et al. 2024) proposes to store frames with different noise levels in a queue, and continuously output one clean frame from the head and enqueue Gaussian noise at the tail in each denoising step. However, the model still faces the temporal flickering issue due to insufficient long-range modeling and discrepancy arising from enqueueing Gaussian tail noise.

**Guidance in Denoising.** In the procedure of image/video denoising, various guidance (*e.g.*, visual tokens in attention (Tewel et al. 2024; Jain et al. 2024) and gradient guidance (Meng et al. 2022; Epstein et al. 2023; An et al. 2024; Chen, Laina, and Vedaldi 2024)) has been investigated to control the content generation. For instance, ConsiStory (Tewel et al. 2024) takes the visual tokens from the reference image to facilitate the content alignment across images in a batch, while FreeDoM (Yu et al. 2023) optimizes the denoising procedure with the gradient of the target energy function.

In short, our work exploits a basic T2V model for long video generation without tuning. The proposed Ouroboros-Diffusion contributes by not only studying how the low-frequency component in latent influences the structural consistency, but also how the guidance in video denoising can be better leveraged to achieve frame-level subject consistency.

## 3 Preliminaries: Video Denoising Approach

**Parallel Denoising.** Latent Video Diffusion Models (LVDMs) perform diffusion processes in the latent space

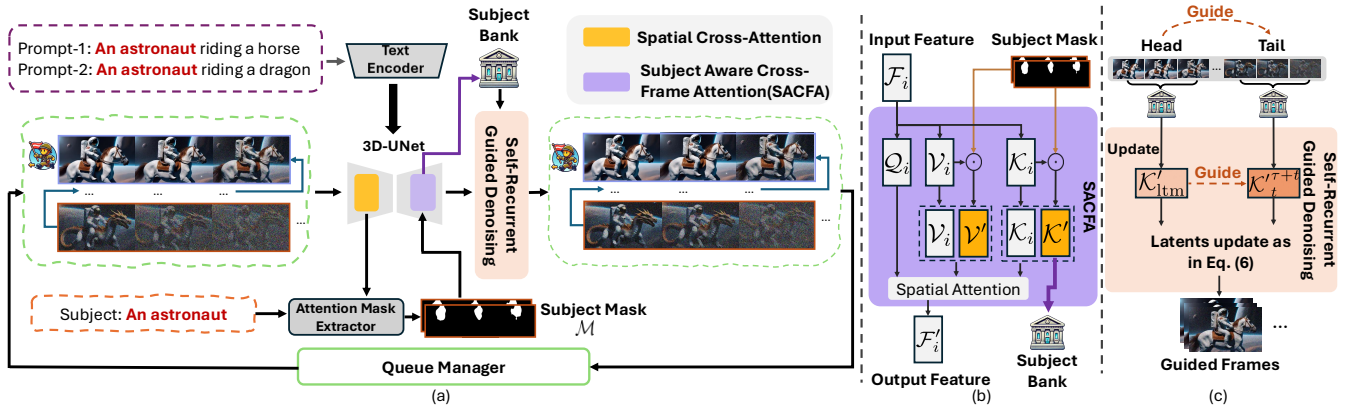


Figure 2: An overview of our Ouroboros-Diffusion. The whole framework (a) contains three key components: coherent tail latent sampling in queue manager , (b) Subject-Aware Cross-frame Attention (SACFA), and (c) self-recurrent guidance. The coherent tail latent sampling in queue manager derives the enqueued frame latents at the queue tail to improve structural consistency. The Subject-Aware Cross-frame Attention (SACFA) aligns subjects across frames within short segments for better visual coherence. The self-recurrent guidance leverages information from all historical cleaner frames to guide the denoising of noisier frames, fostering rich and contextual global information interaction.

for efficient video generation. Given a video latent  $\mathbf{z}_t \in \mathbb{R}^{f \times c \times h \times w}$  at timestep  $t \in [1, \dots, T]$ , where  $\mathbf{z}_t$  consists of  $f$  frame latents:  $\mathbf{z}_t = \{z_t^i\}_{i=1}^f$ ,  $T$  represents the total number of denoising steps used in the sampling process. Conventionally, LVDMs adopt a parallel denoising approach, where  $\mathbf{z}_t$  is iteratively denoised to obtain the clean video latent  $\mathbf{z}_0$ . In parallel denoising, the noise level remains consistent across all frames throughout the denoising process. The parallel denoising is formulated as:

$$\mathbf{z}_{t-1} = \Psi(\mathbf{z}_t, t, \epsilon_\theta(\mathbf{z}_t, t, c)), \quad (1)$$

where  $\Psi(\cdot)$  denotes the sampler such as DDIM and  $\epsilon_\theta$  denotes a spatial-temporal UNet that predicts the added noise at each denoising step, conditioned on the text embedding  $c$ .

**Diagonal Denoising.** Unlike the conventional parallel denoising, FIFO-Diffusion (Kim et al. 2024) introduces a diagonal denoising technique that sequentially produces clear frame latents. This is achieved by employing a fixed-length queue of frame latents with progressively increasing noise levels (*i.e.*, 1 to  $T$ ), as illustrated in Figure 1. In the time step  $\tau$  of diagonal denoising,  $\tau - 1$  frames have already been generated. We denote  $\mathbf{Q}^\tau = \{z_t^{\tau+t}\}_{t=1}^T$  as all frame latents in the queue. Here,  $z_t^{\tau+t}$  is the  $(\tau + t)$ -th frame latent with noise level  $t$ . The denoising step is reformulated as:

$$\{z_{t-1}^{\tau+t}\}_{t=1}^T = \Psi(\mathbf{Q}^\tau, \{t\}_{t=1}^T, \epsilon_\theta(\mathbf{Q}^\tau, \{t\}_{t=1}^T, c)), \quad (2)$$

where  $\{z_{t-1}^{\tau+t}\}_{t=1}^T$  denotes the one-step denoised  $\mathbf{Q}^\tau$ . After the denoising step, the first frame latent  $z_0^\tau$  in the queue becomes a clear latent and is dequeued from the head. A newly sampled Gaussian noise is then enqueued at the tail, making the queue  $\mathbf{Q}^\tau$  to transition to  $\mathbf{Q}^{\tau+1}$ . Iteratively performing this enqueue-dequeue process allows for video generation in a frame-by-frame manner. When the queue size  $T$  exceeds the frame capacity  $f$  of the base video diffusion model, the frame latents are denoised window-by-window with the basic temporal length  $f$ .

## 4 Methodology

Figure 2 provides an overview of the Ouroboros-Diffusion framework. To address the limitations of diagonal denoising, we model three distinct types of information flow to achieve structural and subject consistency. Coherent tail latent sampling ensures smooth transitions by using the second-to-last latent as the guidance, enabling *present* to infer *future*. SACFA enhances subject consistency by extending spatial self-attention with subject contexts from neighboring frames, allowing mutual influence between *present* frames. Self-recurrent guidance further elevates long-range subject coherence by leveraging past subject memory derived from the head of the queue to inform the denoising of the tail, demonstrating how the *past* informs the *present*. In this section, we first discuss the limitations of the FIFO-Diffusion as the motivation of the proposed method. Then we introduce details of our Ouroboros-Diffusion.

### 4.1 Limitation of Diagonal Denoising

The FIFO-Diffusion (Kim et al. 2024) enables the generation of videos with infinite frames. However, the subject consistency of the videos produced by diagonal denoising is often compromised due to the following two limitations:

**Lack of Global Consistency Modeling.** In the design of diagonal denoising, the global visual consistency is not explicitly considered during the diffusion process. Once frames are fully denoised, they are dequeued and no longer contribute to the generation of subsequent frames. This leads to the underutilization of embedded subject information (*e.g.*, semantics, motion). Consequently, later frames cannot reference the information from previously generated frames, which causes the visual appearance of backgrounds and objects to gradually shift during continuous frame denoising, resulting in inconsistencies over time.




Figure 3: The detailed illustration of coherent tail latent sampling in the queue manager.

**Latent Discrepancy Near the Tail.** Another limitation arises in the enqueue-dequeue denoising process, specifically in the construction of the tail latent. The tail latent is typically sampled from a standard Gaussian distribution, which contains pure noise without any visual information. Meanwhile, neighboring latents have already undergone partial denoising in previous steps, introducing some degree of visual content. This creates a discrepancy between the visual information in the tail latent and its neighboring latents, leading to inconsistencies. As a result, the model may struggle to reconcile these differences, often resulting in frame flickering during video generation.

These limitations in diagonal denoising motivate us to devise targeted strategies for improvement.

## 4.2 Coherent Tail Latent Sampling

After completing a DDIM sampling step for queue  $\mathbf{Q}^\tau$ , a clean latent is removed from the queue head, leaving a vacant spot at the tail with a noise level of  $T$ . Instead of sampling Gaussian noise as in Kim et al. (2024), we propose coherent tail latent sampling to retain similar structural information by using the second-to-last latent  $z_{T-1}^{\tau+T}$  as a structural guidance. To ensure the structure similarity between the last two latents, a straightforward approach is to directly apply noise to the second-to-last frame latent and use it as the new tail latent. However, we discover that this approach results in generated videos with limited dynamics due to the excessive similarities in visual content. Recent advances (Wu et al. 2023b; Everaert et al. 2024) indicate that the low-frequency component in the latent space primarily corresponds to the layout and overall structure in the pixel space. As shown in Figure 3, we apply a 2D low-pass filter to extract the low-frequency component of the re-noised latent  $\hat{z}_T^{\tau+T}$ , preserving the layout as the base latent, and introduce dynamics by adding the high-frequency component of a randomly sampled Gaussian noise  $\eta$ . The coherent tail latent sampling is formulated as:

$$z_T^{\tau+1+T} = \mathbf{F}_{\text{low}}^r(z_T^{\tau+T}) + \mathbf{F}_{\text{high}}^r(\eta), \quad (3)$$

where  $\mathbf{F}_{\text{low}}^r(\cdot)$  and  $\mathbf{F}_{\text{high}}^r(\cdot)$  denotes the low-pass and high-pass filter functions with a threshold  $r$ , respectively. The coherent tail latent allows for consistent yet dynamic visual continuation, ensuring a smooth transition from  $\mathbf{Q}^\tau$  to

$\mathbf{Q}^{\tau+1}$ . In this way, the *future* is faithfully inferred from the *present* information within the queue.

## 4.3 Subject-Aware Cross-Frame Attention

To improve subject consistency during denoising, we propose Subject-Aware Cross-Frame Attention (SACFA), which extends the vanilla spatial attention layer to incorporate subject tokens from multiple frames, enhancing visual alignment across frames with enriched subject context.

**Subject Mask Construction.** Central to SACFA is a subject token masking mechanism, which relies on segmenting the subject within each frame. To obtain the segmentation, key subject words are first extracted from a prompt using GPT-4o (OpenAI 2023). These words are then tokenized and encoded into the embedding  $\mathcal{C}_{\text{subj}}$  using the CLIP (Radford et al. 2021) text encoder. The subject tokens are fed into the linear projection layer of each cross-attention layer to form the text subject key  $\mathcal{K}_{\text{subj}}$ . Subject-related attention maps are obtained by computing the attention between the query  $\mathcal{Q}$  and  $\mathcal{K}_{\text{subj}}$  for each cross-attention layer. These maps are averaged across the token dimension and converted into a binary subject mask  $\mathcal{M}$  using Otsu’s method. Finally, subject masks of different resolutions are interpolated to a uniform resolution and averaged, resulting in the final subject mask.

**Attention Processing of SACFA.** To enhance temporal subject consistency across neighboring frames, we extend the spatial self-attention layer into a subject-aware cross-frame approach. The essence is to enable each frame to incorporate subject visual content from other frames when modeling its own appearance. To achieve this, we start by applying subject-specific masks to the keys and values of all frames to extract relevant visual content. These masked keys and values are then concatenated across frames to form subject-aware cross-frame keys and values, denoted as  $\mathcal{K}'$  and  $\mathcal{V}'$ , with the shape  $fh'w' \times d$ . This process creates a collection of subject references that capture the subject information across multiple frames. Finally,  $\mathcal{K}'$  and  $\mathcal{V}'$  are concatenated with the regular key  $\mathcal{K}_i$  and value  $\mathcal{V}_i$  of the target frame  $i$ . The attention processing in SACFA for the  $i$ -th frame is then computed as:

$$\mathcal{F}'_i = \text{Softmax} \left( \frac{\mathcal{Q}_i \cdot [\mathcal{K}_i, \mathcal{K}']^\top}{\sqrt{d}} \right) \cdot [\mathcal{V}_i, \mathcal{V}']. \quad (4)$$

SACFA strengthens local subject correspondence by allowing subject information from neighboring frames to influence current frames, ensuring consistent subject representation. The subject-related key  $\mathcal{K}'$  is then stored in a subject feature bank, forming a subject memory that serves as the foundation for subsequent self-recurrent guidance.

## 4.4 Self-Recurrent Guidance

We introduce self-recurrent guidance to leverage the *past* subject features to guide the *present* denoising steps. Central to this approach is the Subject Feature Bank, which stores the long-term memory of video subjects. The Subject Feature Bank is initialized using the averaged subject-masked keys  $\mathcal{K}'$  from the first  $f$  cleaner frame latents of the video




Figure 4: Visual examples of single-scene long video generation by different approaches. The text prompt is “A cat wearing sunglasses and working as a lifeguard at a pool.”

Approach	Subject Consistency <sup>↑</sup>	Background Consistency <sup>↑</sup>	Motion Smoothness <sup>↑</sup>	Temporal Flickering <sup>↑</sup>	Aesthetic Quality <sup>↑</sup>
StreamingT2V (Henschel et al. 2024)	90.70	95.46	97.34	95.93	54.98
StreamingT2V-VideoTetris (Tian et al. 2024)	89.06	94.80	96.79	95.30	52.89
FIFO-Diffusion (Kim et al. 2024)	94.04	96.08	95.88	93.38	59.06
FreeNoise (Qiu et al. 2024)	94.50	96.45	95.42	93.62	59.32
Ouroboros-Diffusion	<b>96.06</b>	<b>96.90</b>	<b>97.73</b>	<b>96.12</b>	<b>59.89</b>

Table 1: Single-scene video generation performances on VBench. For each video, 128 frames are synthesized for evaluation.

sequence, denoted as  $\mathcal{K}'_{\text{ltm}}$ . These initial  $f$  frames are particularly valuable as they contain clearer and more critical visual information, making them an essential basis for the construction of long-term memory.

After each denoising step of the queue, the bank is updated using an exponential moving average as follows:

$$\mathcal{K}'_{\text{ltm}} \leftarrow \lambda \cdot \mathcal{K}'_{\text{ltm}} + \frac{1-\lambda}{f} \cdot \sum_{t=1}^f \mathcal{K}'_t{}^{\tau+t}, \quad (5)$$

where  $\lambda$  denotes the strength of memorization and only the first  $f$  frames in the queue contribute to the update.

Then, we exploit the subject tokens from the feature bank as a reference to minimize the gradient of the subject discrepancy at the tail. This gradient serves as a guidance to optimize the latent denoising (Yu et al. 2023) as follows:

$$z_{t-1}^{\tau+t} \leftarrow z_{t-1}^{\tau+t} - \gamma_t \cdot \nabla_{z_t^{\tau+t}} \sum_{t=1}^T \|\mathcal{K}'_{\text{ltm}} - \mathcal{K}'_t{}^{\tau+t}\|_2^2, \quad (6)$$

where  $\gamma_t$  denotes a time-dependent strength of the guidance. By integrating this term, the self-recurrent guidance aligns the generated latents more closely with the subject features of the previously generated frames, thereby enhancing long-range subject consistency in the synthesized video.

## 5 Experiments

### 5.1 Experimental Settings

**Benchmark.** We empirically verify the merit of our Ouroboros-Diffusion for both single-scene and multi-scene

long video generation on the VBench (Huang et al. 2024) benchmark. We sample 93 common prompts from VBench as the testing set for single-scene video generation. All methods are required to generate 128 video frames for each prompt. To explore multi-prompt scenarios, we extended the single prompts into multiple prompts using a GPT-4o (OpenAI 2023), resulting in 78 groups of multi-scene prompts. Each group contains 2 to 3 prompts with consistent subject phrasing. For each multi-prompt group, we generate 256 video frames for performance comparison.

**Implementation Details.** We implement our Ouroboros-Diffusion on the text-to-video model VideoCrafter2 (Chen et al. 2024a). The total number of time steps  $T$  in the DDIM sampler is set to 64, matching the queue length. The threshold for the low-pass filter in coherent tail latent sampling is set to 0.25. SACFA is applied only in the down-blocks and mid-block (with down-sampling factors of 2 and 4) of the spatial-temporal UNet empirically. The last 16 frames in the queue are involved in SACFA calculation. The self-recurrent guidance derived from the first 16 frames at the queue head applies to the last 16 frames at the tail. The parameter  $\lambda$  for updating the subject feature bank is set to 0.98.

**Evaluation Metrics.** We choose five evaluation metrics from VBench (Huang et al. 2024) for performance comparison: Subject Consistency, Background Consistency, Motion Smoothness, Temporal Flickering, and Aesthetic Quality. Subject Consistency assesses the uniformity and coherence of the primary subject across frames using DINO (Zhang



Figure 5: Visual examples of multi-scene long video generation by different approaches. The multi-scene prompts are: 1). an astronaut is riding a horse in space; 2). an astronaut is riding a dragon in space; 3). an astronaut is riding a motorcycle in space.

Approach	Subject Consistency $\uparrow$	Background Consistency $\uparrow$	Motion Smoothness $\uparrow$	Temporal Flickering $\uparrow$	Aesthetic Quality $\uparrow$
FIFO-Diffusion (Kim et al. 2024)	93.96	96.17	96.36	93.59	60.12
FreeNoise (Qiu et al. 2024)	95.07	96.52	96.57	95.06	<b>61.26</b>
Ouroboros-Diffusion	<b>95.73</b>	<b>96.82</b>	<b>97.77</b>	<b>95.82</b>	61.17

Table 2: Multi-scene video generation performances on VBench. For each video, 256 frames are synthesized for evaluation.

et al. 2023) features. Background Consistency is measured by the CLIP (Radford et al. 2021) feature similarity. Temporal Flickering evaluates the frame-wise consistency and Motion Smoothness assesses the fluidity and jittering of motion. Finally, Aesthetic Quality indicates the quality of overall visual appearance including composition and color harmony.

## 5.2 Comparisons with State-of-the-Art Methods

We compare our proposal with four state-of-the-art long video diffusion models, *i.e.*, StreamingT2V (Henschel et al. 2024), StreamingT2V-VideoTetris (Tian et al. 2024), FIFO-Diffusion (Kim et al. 2024) and FreeNoise (Qiu et al. 2024) on long video generation.

**Single-Scene Video Generation.** Table 1 summarizes the performance comparison of single-scene long video generation on VBench. Overall, Ouroboros-Diffusion consistently outperforms other baselines across various metrics. Notably, Ouroboros-Diffusion achieves a Temporal Flickering score of 96.12%, surpassing the tuning-free approaches FIFO-Diffusion and FreeNoise by 2.74% and 2.50%, respectively. The highest frame consistency, as indicated by the Temporal Flickering metric, demonstrates the effectiveness of our coherent tail latent sampling, which enforces similarity in image layout between adjacent frames to enhance structural consistency. Additionally, the best performances of Subject Consistency (96.06%) and Background Consistency (96.90%) further show that Ouroboros-Diffusion benefits from subject-level guidance, resulting in natural coherence throughout long video generation. It is important to note that our approach does not compromise video motion strength (*e.g.*, causing static video generation) to improve temporal consistency. To validate this, we calculate the dynamic degree of the pre-trained base model VideoCrafter2 (Chen et al. 2024a), which achieves a score of 42.01. Ouroboros-

Diffusion attains a higher dynamic degree of 44.12, confirming that our method maintains motion variability while further enhancing content consistency in video synthesis.

Figure 4 showcases a single-scene long video generation results across different approaches. Compared to other baselines, Ouroboros-Diffusion consistently produces videos with more seamless transitions and superior visual consistency. For example, StreamingT2V and FreeNoise often generate unreasonable or inconsistent content (*e.g.*, the change of red collar in FreeNoise). Although the FIFO-Diffusion denoising strategy maintains some content coherence (*e.g.*, the appearance of the cat), the enqueued independent Gaussian noise still results in background variations (*e.g.*, changes in the building behind the pool). In contrast, the video generated by our Ouroboros-Diffusion preserves both subject and background consistency effectively, demonstrating the advantage of using information guidance within the queue to enhance visual alignment in diffusion.

**Multi-Scene Video Generation.** Next, we compare our Ouroboros-Diffusion on the task of multi-scene video generation. Table 2 details the performance across different baselines on VBench. Ouroboros-Diffusion outperforms all baselines in Subject/Background Consistency, Motion Smoothness, and Temporal Flickering. Specifically, our approach demonstrates substantial performance boosts (0.66%~1.77%) in Subject Consistency. Note that FreeNoise exploits a noise scheduler for multi-scene video generation, but it emphasizes on prompt adjustment in different denoising steps for motion injection. Our Ouroboros-Diffusion differs fundamentally since ours not only integrates subject visual tokens into cross-frame attention for local alignment but also leverages these moving-averaged tokens to recurrently optimize video latents, ensuring global coherence. Our Aesthetic Quality is slightly lower (by

Model	Coherent Tail Latent Sampling	SACFA	Self-Recurrent Guidance	Subject Consistency $\uparrow$	Background Consistency $\uparrow$	Motion Smoothness $\uparrow$	Temporal Flickering $\uparrow$	Aesthetic Quality $\uparrow$
A	-	-	-	94.04	96.08	95.88	93.38	59.06
B	✓	-	-	95.56	96.66	97.61	95.86	59.61
C	✓	✓	-	95.71	96.73	97.70	96.00	59.67
D	✓	✓	✓	<b>96.06</b>	<b>96.90</b>	<b>97.73</b>	<b>96.12</b>	<b>59.89</b>

Table 3: Performance contribution of each component (i.e., Coherent Tail Latent Sampling, SACFA and Self-Recurrent Guidance) in Ouroboros-Diffusion on single-scene video generation. For each video, 128 frames are synthesized for evaluation.

Model	Motion Smoothness $\uparrow$	Temporal Flickering $\uparrow$
Gaussian Noise	95.88	93.38
Head Frame	97.51	95.87
Second-to-Last Frame (Ours)	<b>97.73</b>	<b>96.12</b>

Table 4: Evaluation on the coherent tail latent sampling.

0.09%) than that of FreeNoise. We speculate that this may be due to a discrepancy between the parallel and diagonal denoising approaches (i.e., consistent noise versus inconsistent noise). This issue could be addressed through model training, and it shows a direction for our future work. Figure 5 further illustrates the multi-scene long video generation results of three different approaches. As shown, Ouroboros-Diffusion successfully generates smoother transitions (e.g., scene changes with the astronaut maintaining the same motion direction) and more consistent visual content (e.g., a single, identical astronaut rather than two).

### 5.3 Ablation Study on Ouroboros-Diffusion

In this section, we conduct ablation studies to evaluate the impact of each design component in Ouroboros-Diffusion for long video generation. All experiments follow previous single-scene video generation settings for comparison.

**Overall Framework.** We first investigate how each component of the overall framework impacts the quality of video generation. Table 3 summarizes the performance results for single-scene long video generation. When integrating coherent tail latent sampling into the base model (A), a significant performance boost (1.52%) is attained by model B in Subject Consistency. This highlights a weakness of the base model (i.e., FIFO-Diffusion), where structural information may be overlooked due to the enqueueing of independent Gaussian noise at the queue tail. By enhancing subject token alignment through subject-aware cross-frame attention, model C outperforms model B across all metrics. Finally, model D, (i.e., our Ouroboros-Diffusion), achieves the best performance by recurrently propagating the subject information of frames from the queue head to the tail frames for latent optimization during video denoising.

**Coherent Tail Latent Sampling.** Next, we present the performance of different variants explored in the design of coherent tail latent sampling. Table 4 details the results of two additional runs: 1) enqueueing independent Gaussian noise at the tail, and 2) replacing the second-to-last frame with the head frame for latent sampling.

Model	Subject Consistency $\uparrow$	Motion Smoothness $\uparrow$
w/o Guidance	94.04	95.88
Guidance with $\lambda=1$	95.87	97.71
Guidance with $\lambda=0$	96.00	97.71
Moving-Average (Ours)	<b>96.06</b>	<b>97.73</b>

Table 5: Evaluation on the self-recurrent guidance.

As expected, independent Gaussian noise yields the lowest Motion Smoothness. Using the head frame’s low-frequency component for latent sampling improves it from 95.88% to 97.51%. However, the appearance gap between queue head and tail latents limits structural guidance. By adjusting enqueued latents with tail information, Ouroboros-Diffusion further enhances motion quality.

**Self-Recurrent Guidance.** We have also analyzed the updating mechanism of the subject feature bank when devising the self-recurrent guidance. As shown in Table 5, whether using historic frame guidance ( $\lambda=1$ ) or the current frame guidance ( $\lambda=0$ ) at the queue head, there are notable performance improvements in Subject Consistency and Motion Smoothness. These results highlight the advantage of guiding the denoising process using subject information through gradient-based adjustments. To achieve better subject consistency in the synthesized video, we implement a moving-average strategy for feature bank updating, which combines information from both historic and current subject tokens. The  $\lambda$  is empirically set as 0.98 in our framework.

## 6 Conclusions

This paper addresses consistent content generation in tuning-free long video diffusion. We introduce Ouroboros-Diffusion, a framework based on the first-in-first-out sampling strategy, which maintains a queue for frame-wise denoising. Our approach examines temporal consistency from two perspectives: structural and subject levels. To materialize this idea, we inject structural information into newly enqueued Gaussian noise by leveraging the low-frequency component of latents near the queue tail, thereby enhancing the consistency of the overall structural layout. At the subject level, we emphasize short-range consistency through visual token alignment in cross-frame attention, while long-range consistency is achieved via latent guidance using subject-aware gradient adjustment. Experiments on the VBench benchmark validate the effectiveness of Ouroboros-Diffusion, demonstrating improvements in both visual quality and motion smoothness.

## References

- An, J.; Yang, Z.; Li, L.; Wang, J.; Lin, K.; Liu, Z.; Wang, L.; and Luo, J. 2024. OpenLEAF: Open-Domain Interleaved Image-Text Generation and Evaluation. In *ACM MM*.
- An, J.; Zhang, S.; Yang, H.; Gupta, S.; Huang, J.-B.; Luo, J.; and Yin, X. 2023. Latent-Shift: Latent Diffusion with Temporal Shift for Efficient Text-to-Video Generation. *arXiv preprint arXiv:2304.08477*.
- Blattmann, A.; Rombach, R.; Ling, H.; Dockhorn, T.; Kim, S. W.; Fidler, S.; and Kreis, K. 2023. Align your Latents: High-Resolution Video Synthesis with Latent Diffusion Models. In *CVPR*.
- Chen, H.; Xia, M.; He, Y.; Zhang, Y.; Cun, X.; Yang, S.; Xing, J.; Liu, Y.; Chen, Q.; Wang, X.; et al. 2023. VideoCrafter1: Open Diffusion Models for High-Quality Video Generation. *arXiv preprint arXiv:2310.19512*.
- Chen, H.; Zhang, Y.; Cun, X.; Xia, M.; Wang, X.; Weng, C.; and Shan, Y. 2024a. VideoCrafter2: Overcoming Data Limitations for High-Quality Video Diffusion Models. *arXiv preprint arXiv:2401.09047*.
- Chen, M.; Laina, I.; and Vedaldi, A. 2024. Training-Free Layout Control with Cross-Attention Guidance. In *WACV*.
- Chen, Z.; Long, F.; Qiu, Z.; Yao, T.; Zhou, W.; Luo, J.; and Mei, T. 2024b. Learning Spatial Adaptation and Temporal Coherence in Diffusion Models for Video Super-Resolution. In *CVPR*.
- Epstein, D.; Jabri, A.; Poole, B.; Efros, A. A.; and Holynski, A. 2023. Diffusion Self-Guidance for Controllable Image Generation. In *NeurIPS*.
- Everaert, M. N.; Fitsios, A.; Bocchio, M.; Arpa, S.; Süssstrunk, S.; and Achanta, R. 2024. Exploiting the Signal-Leak Bias in Diffusion Models. In *WACV*.
- Guo, Y.; Yang, C.; Rao, A.; Liang, Z.; Wang, Y.; Qiao, Y.; Agrawala, M.; Lin, D.; and Dai, B. 2024. AnimateDiff: Animate Your Personalized Text-to-Image Diffusion Models without Specific Tuning. In *ICLR*.
- Henschel, R.; Khachatryan, L.; Hayrapetyan, D.; Poghosyan, H.; Tadevosyan, V.; Wang, Z.; Navasardyan, S.; and Shi, H. 2024. StreamingT2V: Consistent, Dynamic, and Extendable Long Video Generation from Text. *arXiv preprint arXiv:2403.14773*.
- Ho, J.; Chan, W.; Saharia, C.; Whang, J.; Gao, R.; Gritsenko, A.; Kingma, D. P.; Poole, B.; Norouzi, M.; Fleet, D. J.; et al. 2022a. Imagen Video: High Definition Video Generation with Diffusion Models. *arXiv preprint arXiv:2210.02303*.
- Ho, J.; Salimans, T.; Gritsenko, A.; Chan, W.; Norouzi, M.; and Fleet, D. J. 2022b. Video Diffusion Models. In *NeurIPS*.
- Hong, W.; Ding, M.; Zheng, W.; Liu, X.; and Tang, J. 2023. CogVideo: Large-Scale Pretraining for Text-to-Video Generation via Transformers. In *ICLR*.
- Huang, Z.; He, Y.; Yu, J.; Zhang, F.; Si, C.; Jiang, Y.; Zhang, Y.; Wu, T.; Jin, Q.; Chanpaisit, N.; et al. 2024. VBench: Comprehensive Benchmark Suite for Video Generative Models. In *CVPR*.
- Jain, Y.; Nasery, A.; Vineet, V.; and Behl, H. 2024. PEEKA-BOO: Interactive Video Generation via Masked-Diffusion. In *CVPR*.
- Jin, Y.; Sun, Z.; Xu, K.; Chen, L.; Jiang, H.; Huang, Q.; Song, C.; Liu, Y.; Zhang, D.; Song, Y.; Gai, K.; and Mu, Y. 2024. Video-LaVIT: Unified Video-Language Pre-training with Decoupled Visual-Motional Tokenization. *arXiv preprint arXiv:2402.03161*.
- Kim, J.; Kang, J.; Choi, J.; and Han, B. 2024. FIFO-Diffusion: Generating Infinite Videos from Text without Training. *arXiv preprint arXiv:2405.11473*.
- Long, F.; Qiu, Z.; Pan, Y.; Yao, T.; Luo, J.; and Mei, T. 2022a. Stand-Alone Inter-Frame Attention in Video Models. In *CVPR*.
- Long, F.; Qiu, Z.; Pan, Y.; Yao, T.; Ngo, C.-W.; and Mei, T. 2022b. Dynamic Temporal Filtering in Video Models. In *ECCV*.
- Long, F.; Qiu, Z.; Yao, T.; and Mei, T. 2024. VideoStudio: Generating Consistent-Content and Multi-Scene Videos. In *ECCV*.
- Long, F.; Yao, T.; Qiu, Z.; Tian, X.; Luo, J.; and Mei, T. 2019. Gaussian Temporal Awareness Networks for Action Localization. In *CVPR*.
- Long, F.; Yao, T.; Qiu, Z.; Tian, X.; Luo, J.; and Mei, T. 2023. Bi-calibration Networks for Weakly-Supervised Video Representation Learning. *IJCV*.
- Ma, N.; Goldstein, M.; Albergo, M. S.; Boffi, N. M.; Vanden-Eijnden, E.; and Xie, S. 2024a. SiT: Exploring Flow and Diffusion-based Generative Models with Scalable Interpolant Transformers. *arXiv preprint arXiv:2401.08740*.
- Ma, Y.; He, Y.; Cun, X.; Wang, X.; Chen, S.; Li, X.; and Chen, Q. 2024b. Follow Your Pose: Pose-Guided Text-to-Video Generation Using Pose-Free Videos. In *AAAI*.
- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2022. SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations. In *ICLR*.
- Oh, G.; Jeong, J.; Kim, S.; Byeon, W.; Kim, J.; Kim, S.; Kwon, H.; and Kim, S. 2024. MEVG: Multi-event Video Generation with Text-to-Video Models. In *ECCV*.
- OpenAI. 2023. GPT-4 Technical Report. *arXiv:2303.08774*.
- Peebles, W.; and Xie, S. 2023. Scalable Diffusion Models with Transformers. In *ICCV*.
- Peng, B.; Chen, X.; Wang, Y.; Lu, C.; and Qiao, Y. 2024. ConditionVideo: Training-Free Condition-Guided Video Generation. In *AAAI*.
- Qiu, H.; Xia, M.; Zhang, Y.; He, Y.; Wang, X.; Shan, Y.; and Liu, Z. 2024. FreeNoise: Tuning-Free Longer Video Diffusion via Noise Rescheduling. In *ICLR*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning Transferable Visual Models from Natural Language Supervision. In *ICML*.
- Singer, U.; Polyak, A.; Hayes, T.; Yin, X.; An, J.; Zhang, S.; Hu, Q.; Yang, H.; Ashual, O.; Gafni, O.; et al. 2023. Make-A-Video: Text-to-Video Generation without Text-Video Data. In *ICLR*.

Tan, Z.; Yang, X.; Liu, S.; and Wang, X. 2024. Video-Infinity: Distributed Long Video Generation. *arXiv preprint arXiv:2406.16260*.

Tewel, Y.; Kaduri, O.; Gal, R.; Kasten, Y.; Wolf, L.; Chechik, G.; and Atzmon, Y. 2024. Training-free Consistent Text-to-Image Generation. *ACM Transactions on Graphics (TOG)*.

Tian, Y.; Yang, L.; Yang, H.; Gao, Y.; Deng, Y.; Chen, J.; Wang, X.; Yu, Z.; Tao, X.; Wan, P.; et al. 2024. VideoTetris: Towards Compositional Text-to-Video Generation. *arXiv preprint arXiv:2406.04277*.

Villegas, R.; Babaeizadeh, M.; Kindermans, P.-J.; Moraldo, H.; Zhang, H.; Saffar, M. T.; Castro, S.; Kunze, J.; and Erhan, D. 2023. Phenaki: Variable Length Video Generation from Open Domain Textual Description. In *ICLR*.

Voleti, V.; Jolicoeur-Martineau, A.; and Pal, C. 2022. MCVD-Masked Conditional Video Diffusion for Prediction, Generation, and Interpolation. In *NeurIPS*.

Wang, F.-Y.; Chen, W.; Song, G.; Ye, H.-J.; Liu, Y.; and Li, H. 2023a. Gen-L-Video: Multi-Text to Long Video Generation via Temporal Co-Denoising. *arXiv preprint arXiv:2305.18264*.

Wang, X.; Yuan, H.; Zhang, S.; Chen, D.; Wang, J.; Zhang, Y.; Shen, Y.; Zhao, D.; and Zhou, J. 2023b. VideoComposer: Compositional Video Synthesis with Motion Controllability. In *NeurIPS*.

Wu, J. Z.; Ge, Y.; Wang, X.; Lei, W.; Gu, Y.; Shi, Y.; Hsu, W.; Shan, Y.; Qie, X.; and Shou, M. Z. 2023a. Tune-A-Video: One-Shot Tuning of Image Diffusion Models for Text-to-Video Generation. In *ICCV*.

Wu, T.; Si, C.; Jiang, Y.; Huang, Z.; and Liu, Z. 2023b. FreeInit : Bridging Initialization Gap in Video Diffusion Models. *arXiv preprint arXiv:2312.07537*.

Xu, J.; Zou, X.; Huang, K.; Chen, Y.; Liu, B.; Cheng, M.; Shi, X.; and Huang, J. 2024. EasyAnimate: A High-Performance Long Video Generation Method based on Transformer Architecture. *arXiv preprint arXiv:2405.18991*.

Yin, S.; Wu, C.; Liang, J.; Shi, J.; Li, H.; Ming, G.; and Duan, N. 2023a. DragNUWA: Fine-grained Control in Video Generation by Integrating Text, Image, and Trajectory. *arXiv preprint arXiv:2308.08089*.

Yin, S.; Wu, C.; Yang, H.; Wang, J.; Wang, X.; Ni, M.; Yang, Z.; Li, L.; Liu, S.; Yang, F.; et al. 2023b. NUWA-XL: Diffusion over Diffusion for eXtremely Long Video Generation. In *ACL*.

Yu, J.; Wang, Y.; Zhao, C.; Ghanem, B.; and Zhang, J. 2023. Freedom: Training-Free Energy-Guided Conditional Diffusion Model. In *ICCV*.

Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L. M.; and Shum, H.-Y. 2023. DINO: DETR with Improved DeNoising Anchor Boxes for End-to-End Object Detection. In *ICLR*.

Zhang, Z.; Long, F.; Pan, Y.; Qiu, Z.; Yao, T.; Cao, Y.; and Mei, T. 2024. TRIP: Temporal Residual Learning with Image Noise Prior for Image-to-Video Diffusion Models. In *CVPR*.