This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

AR-Diffusion: Asynchronous Video Generation with Auto-Regressive Diffusion

Mingzhen Sun1,2,*Weining Wang1,*Gen Li3Jiawei Liu3Jiahui Sun1,2Wanquan Feng3Shanshan Lao3Siyu Zhou3Qian He3Jing Liu1,21Institute of Automation, Chinese Academy of Sciences3Bytedance Inc.2University of Chinese Academy of Sciences3Bytedance Inc.

Abstract

The task of video generation requires synthesizing visually realistic and temporally coherent video frames. Existing methods primarily use asynchronous auto-regressive models or synchronous diffusion models to address this challenge. However, asynchronous auto-regressive models often suffer from inconsistencies between training and inference, leading to issues such as error accumulation, while synchronous diffusion models are limited by their reliance on rigid sequence length. To address these issues, we introduce Auto-Regressive Diffusion (AR-Diffusion), a novel model that combines the strengths of auto-regressive and diffusion models for flexible, asynchronous video generation. Specifically, our approach leverages diffusion to gradually corrupt video frames in both training and inference, reducing the discrepancy between these phases. Inspired by auto-regressive generation, we incorporate a non-decreasing constraint on the corruption timesteps of individual frames, ensuring that earlier frames remain clearer than subsequent ones. This setup, together with temporal causal attention, enables flexible generation of videos with varying lengths while preserving temporal coherence. In addition, we design two specialized timestep schedulers: the FoPP scheduler for balanced timestep sampling during training, and the AD scheduler for flexible timestep differences during inference, supporting both synchronous and asynchronous generation. Extensive experiments demonstrate the superiority of our proposed method, which achieves competitive and state-ofthe-art results across four challenging benchmarks.

1. Introduction

Video generation aims to create sequences of frames that are both visually realistic and temporally consistent, ensuring that objects in the video appear vivid and their movements are smooth. Current video generation methods

clean noisy	Synchronous Diffusion	Asynchronous Diffusion	Auto- Regression	AR-Diffusion (Ours)	
Templates of timestep compositions					
Composition Space Size	0(1e3)	Ø(1e48)	Ø(16)	O(1e32)	
Train-infer Consistency			×		
Causal Temporal Correlations	×	×			
Classifier-free Guidance	•		×	S	
Varied Length / /ideo Continuatio	n 🙁				

Figure 1. Different generative models employ different constraints on the timestep compositions and thus exhibit different properties.

[1, 9, 11, 13, 14, 17, 20, 21, 33, 43, 45, 51] can be divided into two categories: synchronous and asynchronous models, depending on how noise is applied during training. Synchronous video generation models apply the same level of noise across all video frames during training, ensuring that each frame has the same signal-noise ratio. Video diffusion models [1, 11, 13, 20, 21, 51] are common examples of synchronous method, which apply a shared noise timestep to all frames. These models have demonstrated promising results, but their reliance on equal noise levels and fixedlength sequences limits their ability to generate videos with varied lengths. To address these limitations, some recent studies have tried a "chunked" approach, where multiple frames are generated simultaneously based on a few preceding ones [2, 6, 19, 38]. While this approach helps to reduce computational complexity, it often suffers from temporal inconsistencies and motion discontinuities due to the limited temporal context.

Asynchronous video generation models, in contrast, allow each frame to be corrupted by different noise levels, making them more adaptable to frames with varying com-

^{*} denotes equal contributions.

¹This research was conducted during internship at bytedance.

²https://github.com/iva-mzsun/AR-Diffusion.

plexities. For example, auto-regressive video generative models [9, 14, 17, 33, 43, 45] generate each frame based on previous clear frames and add noise to future frames, which helps create variable-length videos. However, they often face inconsistencies between training and inference, causing issues like error accumulation [34]. Asynchronous diffusion models [4, 18], which apply independent noise to each frame, offer more flexibility but face challenges in training stability and convergence efficiency [18].

To overcome these challenges, we propose Auto-Regressive Diffusion (AR-Diffusion). First, to make the generation process more efficient, we introduce an Auto-Regressive Video Auto-Encoder (AR-VAE), which compresses video frames into compact, continuous latent features. AR-VAE consists of a time-agnostic video encoder that extracts features without considering the temporal correlations, and a temporal causal video decoder that reconstructs the frames with temporal casual attention to maintain temporal consistency. We then propose the AR-Diffusion framework, which combines the strengths of both auto-regressive and diffusion methods. By applying diffusion to gradually corrupt video frames during both training and inference, it reduces discrepancies between these phases, which helps prevent issues like error accumulation. Akin to auto-regressive models, AR-Diffusion employs temporal causal attention and allows earlier video frames to maintain clearer content compared to later ones through asynchronous frame-wise timesteps, thus supporting adaptive inference and facilitating the generation of videos of varied lengths. Unlike previous methods, AR-Diffusion applies a unique non-decreasing constraint on the noise levels: earlier frames remain clearer while later frames get more noise. This helps maintain a natural and coherent progression of generated content, resulting in smoother videos. As illustrated in Fig. 1, compared to existing asynchronous video diffusion models, our nondecreasing constraint can significantly reduce the search space of timestep compositions, thus improving training stability and speeding up convergence.

In addition, we introduce two specialized timestep schedulers. During training, we employ the Frame-oriented Probability Propagation (FoPP) scheduler to balance the sampling of timestep compositions and frame-specific timesteps, ensuring that the model can be well generalized to different inference settings. During inference, we introduce the Adaptive-Difference (AD) scheduler, which allows the timestep difference between neighboring frames to vary adaptively. This flexibility supports both auto-regressive asynchronous and diffusion-based synchronous generation, enhancing adaptability during the inference process.

We have conducted comprehensive experiments to validate the effectiveness and efficiency of our proposed method. The results shows that our method achieves competitive and state-of-the-art (SOTA) performance on four challenging benchmarks, including FaceForensics [26], Sky-Timelapse [42], Taichi-HD [29], and UCF-101 [32]³. In particular, our AR-Diffusion surpasses the previous SOTA asynchronous video diffusion model [18] by 60.1% FVD score on the UCF-101 dataset. Furthermore, we explore the impact of varying timestep differences in video generation across multiple benchmarks, demonstrating that an appropriate timestep difference can significantly enhance the generation performance. Specifically, on the Taichi-HD dataset, an optimal timestep difference can bring a 14.6 and 5.4 improvement in the FVD score compared to synchronous and auto-regressive video generation inference settings, respectively.

Our contributions are summarized as follows:

- We introduce AR-VAE and propose a novel video generative model called AR-Diffusion, which combines the benefits of both asynchronous auto-regressive and synchronous diffusion models.
- We introduce a novel FoPP timestep scheduler during training, which balances uniform sampling of timestep compositions and frame-specific timesteps.
- We propose an AD video scheduler for inference, which allows the timestep difference between neighbor frames to be adaptive.
- Extensive experiments demonstrate the effectiveness and efficiency of our proposed method, validating the necessity of allowing asynchronous timesteps and the importance of finding an optimal timestep difference.

2. Related Works

2.1. Synchronous Video Generation

Synchronous video generation models apply consistent noise or transformations across all video frames, maintaining uniform information entropy and temporal coherence throughout the sequence. Diffusion-based video generative models [1, 5, 10-13, 20, 21, 23, 30, 39, 40, 44, 48-51], are commonly used in this category, as they employ an equal timestep scheduler during both training and inference to ensure consistency across frames. VDM [13] was the first to apply diffusion models to video generation, which introduced a 3D U-Net architecture for video synthesis. Building upon this, models such as LVDM [11] and FDM [10] adopted similar strategies, using identity noise timesteps across different video frames to ensure temporal consistency. Imagen Video [12] and other subsequent works like Make-A-Video [30], Magic Video [51] further improved upon these methods by incorporating spatio-temporally factorized architectures. Recent models such as Lavie [40], PixelDance [49], Open-Sora [50] and Latte [20] have extended these techniques by utilizing larger training datasets and advanced model architectures to enhance generalization and visual fidelity.

³Playable video samples: https://iva-mzsun.github.io/ AR-Diffusion.

In addition to diffusion-based methods, GAN-based approaches have also significantly contributed to synchronous video generation. A collection methods, such as MoCoGAN [36], DIGAN [47], StyleVideoGAN [8], MoStGAN-V [28] and LongVideoGAN [3], achieve temporal consistency by synchronizing content representations across frames while allowing motion-specific variations. Although these methods have shown promising results in capturing temporal dynamics, GAN-based models face challenges such as mode collapse and training instability, which makes it difficult to maintain high-fidelity performance on challenging datasets.

To alleviate the limitations of the equal timestep scheduler, recent researchers have investigated chunked autoregressive generation based on diffusion models, which predict multiple frames in parallel based on a few preceding ones. Methods like MCVD [38], SEINE [6], and VideoFusion [19] reduce computational complexity by generating frames in chunks. However, the temporal context in this approach is limited, thus resulting in inconsistent temporal dynamics.

2.2. Asynchronous Video Generation

Asynchronous video generation methods allow each video frame to be processed with distinct noise levels, offering greater adaptivity in handling varying information entropy across frames. These methods are particularly well-suited for generating videos with variable lengths, as they can be naturally extended to longer tokens or generate subsequent video frames in a first-in-first-out manner[15, 16]. Auto-regressive video generative methods is a representative asynchronous approach [9, 14, 17, 33, 43, 45], which generate each frame sequentially by conditioning on the previously generated ones. This approach allows for flexible, variable-length generation but often introduces inconsistencies between training and inference, leading to issues such as error accumulation and degraded quality in long sequences [34].

Asynchronous diffusion models offer a more flexible approach by applying independent noise timesteps to each video frame[18, 27, 41]. For example, Rolling Diffusion [27] employs a sliding window denoising process for better temporal dynamics. Diffusion Forcing [4] and FVDM [18] use frame-specific noise levels, which enhances sampling flexibility and enables better handling of temporal variability. However, these models still face challenges in achieving stable training [18] and are often outperformed by synchronous methods due to the expanded search space introduced by the independent noise timesteps for each frame.

3. Methods

In this section, we provide a detailed explanation of our proposed method. An overview of the entire framework is illustrated in Fig. 2. We represent a video as $\mathbf{x} \in R^{F \times H \times W \times C}$, where F is the number of video frames, H is the height, W is the width, and C is the number of channels. Our AR-VAE

encodes a video **x** into features $\mathbf{z} \in R^{F \times L \times D}$, where L and D are the length and dimension of latent tokens, respectively. z_i represents the latent feature of the *i*-th video frame x_i . We employ T diffusion timesteps, and the noise timestep for the *i*-th video frame is denoted as t_i . A possible timestep composition for all F video frames can be represented as $\langle t_1, t_2, ... t_F \rangle$, $1 \le t_i \le T$, $1 \le i \le F$. Different diffusion models employ different constraints on these timestep compositions, which will be detailed in Sec. 3.2.

3.1. AR-VAE

Our AR-VAE model is built upon on a Transformer-based 1-Dimensional Tokenizer (Titok) [46], which encodes an image into L discrete tokens. TiTok provides a more compact latent representation, resulting in substantially more efficient and effective encoding compared to conventional techniques. Therefore, our AR-VAE extends Titok's framework to video sequences. As illustrated in Fig. 2(a), we first divide each video frame into frame patches. These frame patches are then concatenated with L learnable video token embeddings into the Transformer-based video encoder. Through crossattention, these token embeddings can effectively capture the video content, thereby serving as visual tokens that represent each video frame. In the original Titok, a vector-quantization module was used to quantize and discretize the image tokens. However, in AR-VAE, we remove this module and instead use the features output by the first normalization layer of the video decoder to represent videos, allowing continuous representation of video features.

During decoding, learnable patch tokens are used as placeholders for the patches of each video frame, and these tokens are shared across different frames. This mechanism ensures that each patch token can interact with others within the same frame to maintain consistency during reconstruction. As shown in the gray solid lines in the video decoder in Fig. 2(a), each patch token is capable of interacting with other patch tokens and visual tokens of the current frame via full attention. Here, we further modify the video decoder to be temporally causal, enhancing temporal coherence and facilitating potential future applications in image auto-encoding. Specifically, patch tokens from subsequent video frames can refer to latent tokens from previous frames, which strengthens their temporal correlations and improves temporal consistency. We represent such reference with dotted and directed lines in Fig. 2(a), indicating that the reference is unidirectional; In other words, video tokens in preceding frames cannot access patch tokens from subsequent video frames. Finally, the video encoder and decoder are together optimized to reconstruct the frame patches. To enhance the clarity of reconstructed video frames, we incorporate an additional adversarial training loss [7]. Moreover, the dimensionality of the video features D is reduced to facilitate the optimization of the follow-up AR-Diffusion by



(a) AR-VAE

(b) AR-Diffusion

Figure 2. The overall framework of our proposed (a) AR-VAE and (b) AR-Diffusion. AR-VAE encodes videos into latent video tokens and AR-Diffusion models asynchronous video generation in the latent space.

removing parameters associated with later dimensions.

3.2. AR-Diffusion

As illustrated in Fig. 2, AR-Diffusion employs a Transformer with temporal causal attention as its backbone. The AR-Diffusion model takes noisy tokens as input and is optimized to output clean tokens. During training, clean video tokens are corrupted according to a sampled timestep composition, and x_0 prediction loss is used to optimize AR-Diffusion. During inference, AR-Diffusion transforms a pure noise into a realistic video sample by iteratively performing the following steps: 1) Predict a clean sample based on the noisy input; 2) Recorrupt the predicted clean sample based on the corresponding timestep composition; 3) Feed the corrupted result into the model for next-step prediction. We discuss the necessity of temporal casual attention and x_0 prediction loss in the appendix.

Diffusion Theory Following [25], each frame feature z_i^0 (i.e. z_i) is corrupted by T steps during the forward diffusion process using the transition kernel:

$$q(z_i^{t_i}|z_i^{t_i-1}) = \mathcal{N}(z_i^{t_i}; \sqrt{1 - \beta_{t_i}} z_i^{t_i-1}, \beta_{t_i} \mathbf{I})$$
(1)

$$q(z_i^{t_i}|z_i^0) = \mathcal{N}(z_i^{t_i}; \sqrt{\bar{\alpha}_{t_i}} z_i^0, (1 - \bar{\alpha}_{t_i})\mathbf{I})$$
(2)

where $\{\beta_t \in (0, 1)\}_{t=1}^T$ is a set of hyper-parameters, $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$. Based on Eq. (2), we can obtain the corrupted feature $z_i^{t_i}$ directly given the timestep t_i during the training and inference phases as follows:

$$z_i^{t_i} = \sqrt{\bar{\alpha}_{t_i}} z_i^0 + (1 - \bar{\alpha}_{t_i}) \epsilon_{t_i}$$
(3)

where ϵ_{t_i} is a noise feature sampled from an isotropic Gauss distribution $\mathcal{N}(0, \mathbf{I})$. The reverse diffusion process $q(z_i^{t_i-1}|z_i^{t_i}, z_i^0)$ has a traceable distribution:

$$q(z_i^{t_i-1}|z_i^{t_i}, z_i^0) = \mathcal{N}(z_i^{t_i-1}|\tilde{\mu}_{t_i}(z_i^{t_i}, z_i^0), \tilde{\beta}_{t_i}\mathbf{I})$$
(4)

where
$$\tilde{\mu}_{t_i}(z_i^{t_i}, z_i^0) = \frac{1}{\sqrt{\alpha_{t_i}}}(z_i^{t_i} - \frac{\beta_{t_i}}{\sqrt{1 - \bar{\alpha}_{t_i}}}\epsilon_{t_i}), \epsilon_{t_i} \sim \mathcal{N}(0, \mathbf{I}),$$

and $\tilde{\beta}_{t_i} = \frac{1 - \bar{\alpha}_{t_i-1}}{1 - \bar{\alpha}_{t_i}}\beta_{t_i}.$

During training, x_0 prediction optimizes the model to predict z_i^0 based on $z_i^{t_i}$. During inference, we can obtain $z_i^{t_i-1}$ based on input $z_i^{t_i}$ and predicted z_i^0 using Eq. 4.

Non-decreasing Timestep Constraint. In synchronous diffusion models [13], a shared noise timestep is applied to all video frames during each training and inference step. This constraint can be formulated as $t_1 = t_2 = ... = t_F$. In a typical scenario with F = 16 and T = 1000, the number of possible timestep compositions satisfying this constraint is $\mathcal{O}(1e3)$, which significantly limits the flexibility of the model. Conversely, recent asynchronous diffusion models [4] allow each video frame to have an independently sampled noise timestep, expanding the number of possible timestep compositions to $\mathcal{O}(1e48)$. However, this vast search space introduces redundancy, as many timestep combinations are never utilized during inference, leading to training instability [18], as discussed in the appendix.

To overcome these challenges, we propose a nondecreasing timestep constraint, which can be formulated as $t_1 \leq t_2 \leq ... \leq t_F$. Intuitively, this constraint ensures that earlier video frames remains as clear or clearer than subsequent frames. Under this constraint, the number of valid timestep compositions in the above-mentioned typical scenario becomes O(1e32), providing greater diversity than the equal constraint while being significantly more stable than the independent sampling approach.

3.3. FoPP Timestep Scheduler

A significant challenge in training our AR-Diffusion is the demand for an appropriate timestep scheduler. During training, the timestep scheduler determines the specific timestep composition utilized to corrupt video frames. It is crucial since it determines how generative models can be general-

Timestep Constraint	Equal	Independent	Non-decreasing
Paths of <u>ALL</u> Timestep compositions			
Composition Space Size	3	27	10
Num of Paths Passed Though Frame-specific Timestep	:1 :1 :1 :1 :1 :1 :1 :1 :1	9: 9 9: 9 9: 9 9: 9 9: 9 9: 9 9: 9 9: 9 9: 9	$\begin{array}{c} :6 : 3 : 1 \\ :3 : 4 : 3 \\ :1 : 3 : 3 : 6 \end{array}$
Uniform Sampling of Timestep Composition vs Frame-specific Timestep	=	=	*

Figure 3. Settings and properties of different timestep constraints when the number of frames and timesteps is set to be F = 3 and T = 3, respectively.

ized to different inference settings. Existing synchronous and asynchronous diffusion models employ timestep schedulers that either equally or independently sample timesteps for video frames. Specifically, synchronous diffusion models adopt an equal timestep scheduler, which first selects a random t from the range of 1 to T and then applies this tas the noise timestep across all video frames. Conversely, asynchronous diffusion models implement a frame-wise independent timestep scheduler. This involves independently sampling a noise timestep $t_i \sim \mathcal{U}(1,T)$ for each video frame sequentially, where i ranges from 1 to F. Both equal and independent timestep schedulers naturally enable both uniformly sampling of timestep compositions and uniformly sampling of frame-specific timesteps, as illustrated in Fig. 3. However, the imposition of the non-decreasing constraint brings conflict between these two sampling uniformity in our AR-Diffusion, as it leads to varying frequencies of framespecific timesteps visited by all possible timestep compositions. To this end, a novel timestep scheduler is required that achieves a balance between these two sampling uniformity when training our AR-Diffusion.

Given previous timestep schedulers, we intuitively design our timestep scheduler by uniformly sampling the noise timestep $t_1 \sim \mathcal{U}(1,T)$ for the first video frame, followed by $t_2 \sim \mathcal{U}(t_1,T)$ for the second frame, and continuing in this manner for subsequent frames. However, this approach results in significantly unbalanced probabilities for different timestep compositions. For instance, if the timestep for the first frame is sampled at T, which occurs with a probability of $\frac{1}{1000}$, then all subsequent frame timesteps must also be T, due to our non-decreasing timestep constraint. Consequently, the composition $\langle t_1 = T, t_2 = T, ..., t_F = T \rangle$ is sampled with a probability of $\frac{1}{1000}$. Given that there are approximately $\mathcal{O}(1e32)$ possible timestep compositions in the typical setting, this probability is excessively high, leading to severe bias in the training process.

To tackle this issue, we introduce the Frame-oriented Probability Propagation (FoPP) timestep scheduler. Firstly, we uniformly sample a frame index $f \sim U(1, F)$ and a

Algorithm 1 Algorithm of the FoPP timestep scheduler

Input: T (Total timesteps), F (Total frames)
Output: $\langle t_1, t_2,, t_F \rangle$
Initializing matrixes $d^s, d^e \in N^{F\mathbf{x}T}$
Randomly select $f \sim \mathcal{U}(1, F)$ and $t_f \sim \mathcal{U}(1, T)$
1: // Sample timesteps for previous video frames
2: for $i = f - 1$ to 1 do
$P^{e} = d^{e}[i, 1: t_{i+1} + 1] / sum(d^{e}[i, 1: t_{i+1} + 1])$
Sample t_f from $\{1,, t_{i+1}\}$ based on probability P^e
3: end for
4: // Sample timesteps for subsequent video frames
5: for $i = f + 1$ to <i>F</i> do
$P^{s} = d^{s}[i, t_{i-1}: T+1] / sum(d^{s}[i, t_{i-1}: T+1])$
Sample t_f from $\{t_{i-1},, T\}$ based on P^s
6: end for

corresponding timestep $t \sim U(1,T)$, ensuring a uniform distribution of timesteps across all video frames.

Then, we utilize the methodology of dynamic programming to calculate and propagate the probability of each timestep for previous and subsequent video frames that can be visited by all possible timestep compositions. In this way, we uniformly sample a timestep composition conditioned on $t_f = t$. In particular, we first use $d_{i,j}^s, 1 \leq i \leq t$ $F, 1 \leq j \leq T$ to represent the total number of timestep compositions that satisfies $\langle t_i = j, t_{i+1}, ..., t_F \rangle$ and the nondecreasing constraint. Here, $\langle t_i = j, t_{i+1}, ..., t_F \rangle$ means that each timestep composition is starting with the *i*-th video frame at timestep $t_i = j$. Then, we calculate $d_{i,j}^s$ with i from F to 1 and j from T to 1 using dynamic programming, where $d_{*,T} = 1$, $d_{F,*} = 1$, and the transition equation being $d_{i,j} = d_{i,j-1} + d_{i-1,j}$. Then the visit probability of the timestep $k \in [K, T]$ of the subsequent video frame *i*, given i > f and $t_{i-1} = K$, is $\frac{d_{i,k}^s}{\sum_{j=K}^T d_{i,j}^s}$. Similarly, we calculate $d_{i,j}^e$, which represents the total number of timestep compositions that satisfies $\langle t_1, t_2, ..., t_i = j \rangle$ and the non-decreasing constraint. Then the visit probability of the timestep $k \in [1, K]$ of the previous video frame *i*, given $i < f \text{ and } t_{i+1} = K, \text{ is } \frac{d_{i,k}^e}{\sum_{j=1}^K d_{i,j}^e}.$

Finally, the timesteps of previous or subsequent video frames are sampled one by one based on corresponding probabilities. The methodology is detailed in Algorithm 1.

3.4. AD Timestep Scheduler

During inference, a timestep scheduler is utilized to regulate a sequence of timestep compositions that starts from $\langle t_1 = T, t2 = T, ..., t_F = T \rangle$ to $\langle t_1 = 0, t2 = 0, ..., t_F = 0 \rangle$. In other words, it determines which timestep composition is used in each sampling step. Here, we introduce an AD timestep scheduler, which supports adaptive video generation and accommodates both asynchronous auto-regressive and synchronous generation. Specifically, the AD timestep scheduler makes the timestep difference between neighboring video frames as an adaptive variable s. For timesteps t_i and t_{i-1} of consecutive video frames, we ensure that the following condition holds:

$$t_{i} = \begin{cases} t_{i} + 1, & \text{if } i = 1 \text{ or } t_{i-1} = 0, \\ min(t_{i-1} + s, T), & \text{if } t_{i-1} > 0 \end{cases}$$
(5)

Specifically, when the previous video frame is none or clean, the current video frame focuses on the denoising of itself; Otherwise, it keeps denoising with a timestep difference of s than its previous video frame. Notably, the synchronous diffusion model and the auto-regressive generative model represent two distinct cases within our variable difference timestep scheduler, corresponding to s = 0 and s = T, respectively. Intuitively, a smaller s results in neighboring video frames having more similar content, whereas a larger s introduces greater content variability. Importantly, we apply a mask to those frame latents for which the noise timesteps remain unchanged.

4. Experiments

4.1. Setups

Datasets. Following previous video generation methods [33, 34], we conduct our experiments on four challenging datasets: FaceForensics [26], Sky-Timelapse [42], Taichi-HD [29], and UCF-101 [32].

Evaluation metrics. For quantitative evaluation, we use FID-img, FID-vid, and FVD [37] to measure the realism of generated videos. For all metrics, the lower the score, the better the effect. We consider two different clip lengths (16 and 128) for FVD, where we denote FVD_{16} and FVD_{128} as the FVD score measured on video clips of lengths 16 and 128, respectively.

Implementation Details. Our experiments were conducted on videos with a resolution of 256^2 and on 8 NVIDIA A800-40G GPUs. For the AR-VAE training, we configured the length and dimension of the video latent features at L = 32and D = 4, respectively. We initialized our AR-VAE using the publicly available checkpoint from Titok [46]. Initially, each AR-VAE was trained for 50,000 steps on individual video frames using typical reconstruction loss [24]. Subsequently, the adversarial loss was incorporated for an additional 450,000 training steps to further refine the model. Finally, we integrated temporal causal attention into our video decoder within the AR-VAE and conducted optimization across all parameters over an additional 100,000 steps. The batch size is 16 per GPU and the learning rate is 1e-4. We developed our AR-Diffusion using a Transformer-based diffusion model, specifically DiT [22]. Initially, video latent features are rescaled by a factor of 0.5 before input into AR-Diffusion. The training configuration includes a batch size of 16 per GPU, with gradient accumulation over 8 steps.

The learning rate is set at 2e-4 for the initial 100,000 training steps, followed by a fine-tuning phase for an additional 50,000 steps at a reduced rate of 1e-5. Techniques such as exponential moving average over all parameters, gradient clipping, and noise clipping are also utilized.

4.2. Quantitaive Comparison

We compare our proposed AR-Diffusion with other methods on four challenging datasets, including FaceForensics [26], Sky-Timelapse [42], Taichi-HD [29] and UCF-101 [32]. As shown in Table 1, our AR-Diffusion model achieves lower FVD_{16} and FVD_{128} scores compared to other generative methods, which indicates better video quality and temporal consistency. For example, on the FaceForensics dataset, AR-Diffusion achieves an FVD_{16} of 71.9 and an FVD_{128} of 265.7, respectively, outperforming most of the other models, including PVDM and Diffusion Forcing. On the Sky-Timelapse dataset, AR-Diffusion achieves an FVD_{16} of 40.8, which is lower than all the other methods, showcasing its superior performance in generating high-quality temporal sequences. On the Taichi-HD dataset, AR-Diffusion achieves an FVD_{16} of 66.3 and an FVD_{128} of 376.3, outperforming all the other models including synchronous and asynchronous diffusion-based models, as well as generative adversarial models. On the more complex UCF-101 dataset, AR-Diffusion achieves an FVD_{16} of 186.6 and an FVD_{128} score of 572.3, which significantly outperforms all the other models. These results clearly demonstrate the superiority of AR-Diffusion across diverse datasets, establishing a new benchmark in video generation quality and consistency.

Table 2 presents quantitative results for AR-Diffusion with different settings of the timestep difference s, which controls the generation process and provides a trade-off between generation performance and efficiency. Notably, s = 0 corresponds to synchronous diffusion generation, while s = 50 represents auto-regressive generation. The table reports performance on three metrics: FID-img, FID-vid, and FVD for both 16-frame and 128-frame video generation. For 16-frame video generation, AR-Diffusion generally achieves competitive FVD scores with s values between 0 and 15, depending on the dataset. For example, on the Face-Forensics and Sky-Timelapse datasets, the best FVD scores are achieved when s = 0, indicating that synchronous generation provides optimal quality in this context. On Taichi-HD, s = 10 provides the lowest FVD scores 66.3, suggesting that a moderate combination of synchronous and auto-regressive steps yields improved performance for this dataset. On UCF-101, the lowest FVD score of 186.6 is obtained with s = 15, showing the effectiveness of moderate timestep values in balancing temporal coherence and quality for complex datasets. As s increases beyond 15, FVD scores tend to increase, which suggests a reduction in the temporal quality of generated videos. In addition, inference time also increases with

Table 1. (Quantitative comp	parison on four	challenging dat	asets: Taichi-HD	[29], Sky	-Timelapse [4	42], Fa	ceForensics [2	6], and UCF	-101 [32].
------------	-------------------	-----------------	-----------------	------------------	-----------	---------------	---------	----------------	-------------	------------

	Taichi-	HD [29]	[29] Sky-Timelapse [42]		FaceFore	nsics [26]	UCF-101 [32]		
	FVD ₁₆	FVD ₁₂₈	FVD ₁₆	FVD ₁₂₈	FVD ₁₆	FVD ₁₂₈	FVD_{16}	FVD ₁₂₈	
Generative Adversarial Models									
MoCoGAN [36]	-	-	206.6	575.9	124.7	257.3	2886.9	3679.0	
+ StyleGAN2 backbone	-	-	85.9	272.8	55.6	309.3	1821.4	2311.3	
MoCoGAN-HD [35]	-	-	164.1	878.1	111.8	653.0	1729.6	2606.5	
DIGAN [47]	128.1	748.0	83.11	196.7	62.5	1824.7	1630.2	2293.7	
StyleGAN-V [31]	143.5	691.1	79.5	197.0	47.4	89.3	1431.0	1773.4	
MoStGAN-V [28]	-	-	65.3	162.4	39.7	72.6	-	-	
			Auto-regressive	Generative Mode	els				
VideoGPT [43]	-	-	222.7	-	185.9	-	2880.6	-	
TATS [9]	94.6	-	132.6	-	-	-	420	-	
		Syı	nchronous Diffus	ion Generative N	Iodels				
Latte [20]	159.6	-	59.8	-	34.0	-	478.0	-	
PVDM [48]	540.2	-	55.4	125.2	355.9	-	343.6	648.4	
LVDM [11]	99.0	-	95.2	-	-	-	372.9	-	
VIDM [21]	121.9	563.6	57.4	140.9	-	-	294.7	1531.9	
Asynchronous Diffusion Generative Models									
FVDM [18]	194.6	-	106.1	-	55.0	-	468.2	-	
Diffusion Forcing [4]	202.0	738.5	251.9	895.3	99.5	555.2	274.5	836.3	
AR-Diffusion (ours)	66.3	376.3	40.8	175.5	71.9	265.7	186.6	572.3	

Table 2. Quantitative results with different values of timestep s. We utilize a DDIM sampler with 50 sampling steps in total.

	Face	Forensics [26	5]	Sky-7	Timelapse [4	2]	Taichi-HD [29]		UCF-101 [32]			Inference	
8	FID-img	FID-vid	FVD	FID-img	FID-vid	FVD	FID-img	FID-vid	FVD	FID-img	FID-vid	FVD	Time (s)
16-frame Video Generation													
0	14.0	6.9	71.9	10.0	9.2	40.8	13.8	9.2	80.9	30.3	17.6	194.4	2.4
5	14.0	6.2	78.1	11.1	11.6	55.2	13.0	5.9	66.3	30.0	17.7	194.0	5.2
10	13.6	6.1	84.4	10.3	11.2	57.6	12.4	5.8	70.9	30.1	18.8	212.2	7.9
15	14.3	6.6	83.5	9.4	11.4	55.6	12.2	5.8	69.4	30.0	16.3	186.6	10.8
20	14.8	6.3	83.3	9.2	10.9	56.3	12.7	6.0	67.0	31.0	17.4	201.1	13.6
25	14.1	6.1	79.0	9.7	10.5	48.4	12.9	6.5	75.1	29.6	17.3	191.6	16.4
50	14.2	6.1	82.8	10.1	10.7	50.6	13.1	5.9	71.7	29.5	17.1	192.6	30.5
	128-frame Video Generation												
5	14.7	9.5	265.7	12.2	25.2	185.1	8.9	10.8	376.3	32.5	24.4	592.7	42.1
10	15.2	8.9	278.1	12.1	23.9	182.6	8.8	12.2	401.9	31.5	24.9	605.3	78.0
25	15.4	9.3	348.6	12.2	22.8	175.5	8.8	12.3	402.5	31.8	23.3	572.3	184.8

larger s values, indicating a trade-off between efficiency and the use of auto-regressive steps. For 128-frame video generation, s = 5 achieves relatively lower FVD scores, demonstrating that a smaller s value is beneficial for generating longer sequences while maintaining visual quality. Overall, the analysis indicates that a moderate value of s (e.g., s = 10 for 16-frame generation and s = 5 for 128frame generation) often strikes an effective balance between quality and inference efficiency across different datasets.

4.3. Qualitative Comparison

We qualitatively compare our method with prior works as in Figure 4. Samples of prior methods are obtained from [19, 21] or generated with their released codes and parameters [20]. On UCF-101, the GAN-based method DIGAN tends to produce video samples that lack distinctiveness. In contrast, TATS, which leverages a Transformer architecture and an interpolation strategy, creates more recognizable video samples. VIDM results in overly simplistic object motions. The asynchronous diffusion model Diffusion Forcing [4] yields videos with minimal motion, while our AR-Diffusion method produces samples that showcase both clear appearances and notable movements. On the Sky Time-lapse dataset, samples from DIGAN and TATS exhibit minimal motion and simplistic object representations. Conversely, Latte and VIDM enhance the detail and clarity of object boundaries, though their motion effects are still limited. In contrast, our AR-Diffusion method creates video samples characterized by dynamic movements and significantly enhanced visual details. Similarly, for the TaiChi-HD dataset, video samples from DIGAN and TATS show noticeable distortions in human appearances. While Latte, VIDM, and Diffusion Forcing [4] improve the depiction of human figures, their movements remain minimal. In contrast, samples produced by our AR-Diffusion method display significant movements and clearly defined object appearances.

4.4. Ablation Study

To systematically evaluate the contributions of different model components in our proposed AR-Diffusion model, we design various ablation models as shown in Table 3. Each ablation experiment is conducted on the Sky-timelapse dataset using 8 A800 GPUs without the fine-tuning stage. The baseline AR-Diffusion model achieves an FID of 12.2, an FVD-



Figure 4. Qualitative comparison of existing video generative methods and our AR-Diffusion.

Table 3. Ablation study on our AR-Diffusion.

	FID	FVD-img	FVD
AR-Diffusion	12.2	13.4	62.8
-FoPP Timestep Scheduler	11.0	16.8	101.0
-Improved VAE	13.1	29.6	148.3
-Temporal Causal Attention	15.9	50.2	209.8
-x0 Prediction Loss	27.9	58.0	257.6
-Non-decreasing Constraint	32.2	87.9	272.5

img of 13.4 and an FVD of 62.8, indicating high video quality and temporal consistency. Removing the FoPP timestep scheduler leads to an increase in FVD to 101.0, demonstrating the necessity of appropriately sampling timestep composition during training. Our AR-VAE is improved by reducing the latent dimension from 12 to 4, employing temporal causal attention, and incorporating the adversarial training loss. Removing these improvements, the FVD score further increases to 148.3, illustrating their significant impact on generating high-quality video content. Eliminating temporal causal attention in the AR-Diffusion causes a substantial performance drop, with FID increasing to 15.9 and FVD rising to 209.8, underscoring its importance in maintaining temporal dependencies between frames. Replacing x_0 prediction with ϵ prediction has a particularly drastic effect, with FID increasing to 27.9 and FVD reaching 257.6, suggesting that this loss function is crucial for stable and consistent video generation. Finally, the absence of the non-decreasing constraint leads to the worst performance, with FID-img and FID-vid increasing to 32.2 and 87.9, respectively, and FVD reaching 272.5, indicating that

this constraint is essential for maintaining quality during diffusion steps. Overall, the ablation study demonstrates that each component significantly contributes to the performance of AR-Diffusion, particularly the temporal causal attention, improved VAE, and x_0 prediction loss.

5. Conclusion

In this paper, we present a novel auto-regressive video diffusion model, which is called AR-Diffusion. Targeting asynchronous video generation, AR-Diffusion introduces a novel non-decreasing timestep constraint, which significantly reduces the search space and thus stabilizes the training procedure. During training, the FoPP timestep scheduler is used to optimize AR-Diffusion with a balanced sampling of timestep compositions. During inference, the AD timestep scheduler is employed to regulate a sequence of timestep compositions for AR-Diffusion to produce a video sample. In conclusion, AR-Diffusion combines the advantages of both auto-regressive generative models and synchronous video diffusion models, obtaining competitive or SOTA results on four challenging benchmarks.

Acknowledgement. This research is supported by Artificial Intelligence-National Science and Technology Major Project 2023ZD0121200and the National Natural Science Foundation of China (6243000159 62102416), and the Key Research and Development Program of Jiangsu Province under Grant BE2023016-3.

References

- [1] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127, 2023. 1, 2
- [2] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 22563–22575, 2023. 1
- [3] Tim Brooks, Janne Hellsten, Miika Aittala, Ting-Chun Wang, Timo Aila, Jaakko Lehtinen, Ming-Yu Liu, Alexei Efros, and Tero Karras. Generating long videos of dynamic scenes. Advances in Neural Information Processing Systems, 35:31769– 31781, 2022. 3
- [4] Boyuan Chen, Diego Marti Monso, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. Diffusion forcing: Next-token prediction meets full-sequence diffusion. *arXiv preprint arXiv:2407.01392*, 2024. 2, 3, 4, 7
- [5] Haoxin Chen, Menghan Xia, Yingqing He, Yong Zhang, Xiaodong Cun, Shaoshu Yang, Jinbo Xing, Yaofang Liu, Qifeng Chen, Xintao Wang, et al. Videocrafter1: Open diffusion models for high-quality video generation. arXiv preprint arXiv:2310.19512, 2023. 2
- [6] Xinyuan Chen, Yaohui Wang, Lingjun Zhang, Shaobin Zhuang, Xin Ma, Jiashuo Yu, Yali Wang, Dahua Lin, Yu Qiao, and Ziwei Liu. Seine: Short-to-long video diffusion model for generative transition and prediction. In *International Conference on Learning Representations*, 2023. 1, 3
- [7] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *IEEE* conference on computer vision and pattern recognition, pages 12873–12883, 2021. 3
- [8] Gereon Fox, Ayush Tewari, Mohamed Elgharib, and Christian Theobalt. Stylevideogan: A temporal generative model using a pretrained stylegan. In *British Machine Vision Conference*, 2021. 3
- [9] Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and Devi Parikh. Long video generation with time-agnostic vqgan and time-sensitive transformer. In *European Conference on Computer Vision*, pages 102–118, 2022. 1, 2, 3, 7
- [10] William Harvey, Saeid Naderiparizi, Vaden Masrani, Christian Weilbach, and Frank Wood. Flexible diffusion modeling of long videos. Advances in Neural Information Processing Systems, 35:27953–27965, 2022. 2
- [11] Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. Latent video diffusion models for high-fidelity long video generation. arXiv preprint arXiv:2211.13221, 2022. 1, 2, 7
- [12] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. arXiv preprint arXiv:2210.02303, 2022. 2

- [13] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. Advances in Neural Information Processing Systems, 35:8633–8646, 2022. 1, 2, 4
- [14] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *International Conference on Learning Representations*, 2022. 1, 2, 3
- [15] Jihwan Kim, Junoh Kang, Jinyoung Choi, and Bohyung Han. Fifo-diffusion: Generating infinite videos from text without training. arXiv preprint arXiv:2405.11473, 2024. 3
- [16] Akio Kodaira, Chenfeng Xu, Toshiki Hazama, Takanori Yoshimoto, Kohei Ohno, Shogo Mitsuhori, Soichi Sugano, Hanying Cho, Zhijian Liu, and Kurt Keutzer. Streamdiffusion: A pipeline-level solution for real-time interactive generation. arXiv preprint arXiv:2312.12491, 2023. 3
- [17] Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Grant Schindler, Rachel Hornung, Vighnesh Birodkar, Jimmy Yan, Ming-Chang Chiu, et al. Videopoet: A large language model for zero-shot video generation. *arXiv preprint arXiv:2312.14125*, 2023. 1, 2, 3
- [18] Yaofang Liu, Yumeng Ren, Xiaodong Cun, Aitor Artola, Yang Liu, Tieyong Zeng, Raymond H Chan, and Jean-michel Morel. Redefining temporal modeling in video diffusion: The vectorized timestep approach. *arXiv preprint arXiv:2410.03160*, 2024. 2, 3, 4, 7
- [19] Zhengxiong Luo, Dayou Chen, Yingya Zhang, Yan Huang, Liang Wang, Yujun Shen, Deli Zhao, Jingren Zhou, and Tieniu Tan. Videofusion: Decomposed diffusion models for high-quality video generation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 10209–10218, 2023. 1, 3, 7
- [20] Xin Ma, Yaohui Wang, Gengyun Jia, Xinyuan Chen, Ziwei Liu, Yuan-Fang Li, Cunjian Chen, and Yu Qiao. Latte: Latent diffusion transformer for video generation. arXiv preprint arXiv:2401.03048, 2024. 1, 2, 7
- [21] Kangfu Mei and Vishal Patel. Vidm: Video implicit diffusion models. In AAAI conference on artificial intelligence, pages 9117–9125, 2023. 1, 2, 7
- [22] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *IEEE International Conference on Computer Vision*, pages 4195–4205, 2023. 6
- [23] PKU-Yuan Lab and Tuzhan AI. Open-sora-plan, 2024. 2
- [24] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831, 2021. 6
- [25] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 4
- [26] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics: A large-scale video dataset for forgery detection in human faces. arXiv preprint arXiv:1803.09179, 2018. 2, 6, 7

- [27] David Ruhe, Jonathan Heek, Tim Salimans, and Emiel Hoogeboom. Rolling diffusion models. *arXiv preprint arXiv:2402.09470*, 2024. 3
- [28] Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Mostganv: Video generation with temporal motion styles. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5652–5661, 2023. 3, 7
- [29] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. Advances in neural information processing systems, 32, 2019. 2, 6, 7
- [30] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, Devi Parikh, Sonal Gupta, and Yaniv Taigman. Makea-video: Text-to-video generation without text-video data. In *International Conference on Learning Representations*, 2023.
- [31] Ivan Skorokhodov, Sergey Tulyakov, and Mohamed Elhoseiny. Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. In *IEEE conference on computer vision and pattern recognition*, pages 3626–3636, 2022. 7
- [32] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402, 2012. 2, 6, 7
- [33] Mingzhen Sun, Weining Wang, Xinxin Zhu, and Jing Liu. Moso: Decomposing motion, scene and object for video prediction. In *IEEE conference on computer vision and pattern recognition*, pages 18727–18737, 2023. 1, 2, 3, 6
- [34] Mingzhen Sun, Weining Wang, Zihan Qin, Jiahui Sun, Sihan Chen, and Jing Liu. Glober: coherent non-autoregressive video generation via global guided video decoder. Advances in Neural Information Processing Systems, 36, 2024. 2, 3, 6
- [35] Yu Tian, Jian Ren, Menglei Chai, Kyle Olszewski, Xi Peng, Dimitris N Metaxas, and Sergey Tulyakov. A good image generator is what you need for high-resolution video synthesis. In *International Conference on Learning Representations*, 2021. 7
- [36] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018. 3, 7
- [37] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. arXiv preprint arXiv:1812.01717, 2018. 6
- [38] Vikram Voleti, Alexia Jolicoeur-Martineau, and Chris Pal. Mcvd-masked conditional video diffusion for prediction, generation, and interpolation. *Advances in neural information processing systems*, 35:23371–23385, 2022. 1, 3
- [39] Weimin Wang, Jiawei Liu, Zhijie Lin, Jiangqiao Yan, Shuo Chen, Chetwin Low, Tuyen Hoang, Jie Wu, Jun Hao Liew, Hanshu Yan, et al. Magicvideo-v2: Multi-stage high-aesthetic video generation. arXiv preprint arXiv:2401.04468, 2024. 2
- [40] Yaohui Wang, Xinyuan Chen, Xin Ma, Shangchen Zhou, Ziqi Huang, Yi Wang, Ceyuan Yang, Yinan He, Jiashuo Yu, Peiqing Yang, et al. Lavie: High-quality video gener-

ation with cascaded latent diffusion models. *arXiv preprint* arXiv:2309.15103, 2023. 2

- [41] Tong Wu, Zhihao Fan, Xiao Liu, Hai-Tao Zheng, Yeyun Gong, Jian Jiao, Juntao Li, Jian Guo, Nan Duan, Weizhu Chen, et al. Ar-diffusion: Auto-regressive diffusion model for text generation. Advances in Neural Information Processing Systems, 36:39957–39974, 2023. 3
- [42] Wei Xiong, Wenhan Luo, Lin Ma, Wei Liu, and Jiebo Luo. Learning to generate time-lapse videos using multi-stage dynamic generative adversarial networks. In *IEEE Conference* on Computer Vision and Pattern Recognition, pages 2364– 2373, 2018. 2, 6, 7
- [43] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation using vq-vae and transformers. arXiv preprint arXiv:2104.10157, 2021. 1, 2, 3, 7
- [44] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. arXiv preprint arXiv:2408.06072, 2024. 2
- [45] Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Vighnesh Birodkar, Agrim Gupta, Xiuye Gu, et al. Language model beats diffusion-tokenizer is key to visual generation. arXiv preprint arXiv:2310.05737, 2023. 1, 2, 3
- [46] Qihang Yu, Mark Weber, Xueqing Deng, Xiaohui Shen, Daniel Cremers, and Liang-Chieh Chen. An image is worth 32 tokens for reconstruction and generation. *arXiv preprint* arXiv:2406.07550, 2024. 3, 6
- [47] Sihyun Yu, Jihoon Tack, Sangwoo Mo, Hyunsu Kim, Junho Kim, Jung-Woo Ha, and Jinwoo Shin. Generating videos with dynamics-aware implicit generative adversarial networks. In *International Conference on Learning Representations*, 2022. 3, 7
- [48] Sihyun Yu, Kihyuk Sohn, Subin Kim, and Jinwoo Shin. Video probabilistic diffusion models in projected latent space. In *IEEE conference on computer vision and pattern recognition*, pages 18456–18466, 2023. 2, 7
- [49] Yan Zeng, Guoqiang Wei, Jiani Zheng, Jiaxin Zou, Yang Wei, Yuchen Zhang, and Hang Li. Make pixels dance: Highdynamic video generation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8850–8860, 2024. 2
- [50] Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-sora: Democratizing efficient video production for all, march 2024. URL https://github. com/hpcaitech/Open-Sora, 2024. 2
- [51] Daquan Zhou, Weimin Wang, Hanshu Yan, Weiwei Lv, Yizhe Zhu, and Jiashi Feng. Magicvideo: Efficient video generation with latent diffusion models. *arXiv preprint arXiv:2211.11018*, 2022. 1, 2