
InfinityStar: Unified Spacetime AutoRegressive Modeling for Visual Generation

Jinlai Liu*, Jian Han*, Bin Yan* Hui Wu, Fengda Zhu, Xing Wang

Yi Jiang, Bingyue Peng, Zehuan Yuan[†]

ByteDance

{liujinlai.licio, hanjian.thu123, bin.yan, wuhui.321, fengdazhu}@bytedance.com,
{xing.wang, jiangyi.enjoy, bingyue.peng, yuanzehuan}@bytedance.com,

Codes and models: <https://github.com/FoundationVision/InfinityStar>

Abstract

We introduce InfinityStar, a unified spacetime autoregressive framework for high-resolution image and dynamic video synthesis. Building on the recent success of autoregressive modeling in both vision and language, our purely discrete approach jointly captures spatial and temporal dependencies within a single architecture. This unified design naturally supports a variety of generation tasks such as text-to-image, text-to-video, image-to-video, and long interactive video synthesis via straightforward temporal autoregression. Extensive experiments demonstrate that InfinityStar scores 83.74 on VBench, outperforming all autoregressive models by large margins, even surpassing some diffusion competitors like HunyuanVideo. Without extra optimizations, our model generates a 5s, 720p video approximately $10\times$ faster than leading diffusion-based methods. To our knowledge, InfinityStar is the first discrete autoregressive video generator capable of producing industrial-level 720p videos. We release all code and models to foster further research in efficient, high-quality video generation.

1 Introduction

Visual synthesis has witnessed remarkable progress in recent years, largely propelled by the scaling of Transformer architectures. In particular, video generation has attracted growing interest from both academia and industry, owing to its wide-ranging applications in content creation, world simulation, etc. At present, diffusion models[3, 20, 19, 32, 9, 47] lead the field by iteratively denoising latent representations to produce high-fidelity clips. Concurrently, autoregressive models[18, 34, 10] have been explored for their potential to unify image and video generation and to generalize over longer time horizons.

Despite their successes, each paradigm exhibits critical shortcomings. Video diffusion models excel at synthesizing fixed-length frame sequences by exploiting bidirectional attention, yet they incur substantial computational cost due to tens or even hundreds of sequential denoising steps, and they struggle to extend seamlessly to video extrapolation. Autoregressive methods based on next-token prediction, while inherently capable of streaming generation, often fall short in visual fidelity and suffer from prohibitive latency due to tens of thousands of inference steps.

These observations motivate the need for a generation framework that simultaneously possess high visual quality, efficiency and temporal generalization. Recently, Visual AutoRegressive modeling (VAR)[29] redefined image generation as a coarse-to-fine next-scale prediction. Its follow-up work, Infinity [15] further introduces bitwise modeling and scales up the vocabulary size, achieving

*Equal contribution. [†]Corresponding author: yuanzehuan@bytedance.com

comparable performance to diffusion models while offering significant advantages in inference speed. Inspired by the success of VAR [29] and Infinity [15], we present InfinityStar, a Spacetime Pyramid Modeling for unified text-to-image, text-to-video, zero-shot image-to-video, and zero-shot video extrapolation. This framework models a video as an image pyramid and multiple clip pyramids, not only naturally inheriting the text-to-image capabilities but also decoupling static appearance from dynamic motions in videos. Furthermore, we introduce several key improvements. First, we improve discrete reconstruction quality by leveraging knowledge inheritance from a continuous video tokenizer. Second, we introduce Stochastic Quantizer Depth during training of the tokenizer to alleviate the imbalanced information distribution across scales. Third, we propose Semantic Scales Repetition, which refines the predictions of early semantic scales in a video, significantly enhancing fine-grained details and complex motions of the generated videos.

We train InfinityStar on large-scale video corpora to support generating videos of up to 720p resolution and variable durations. On the VBench benchmark[45], InfinityStar establishes a new state-of-the-art among autoregressive video generation models, even surpassing industry-leading HunyuanVideo[19] (83.74 v.s 83.24). Besides, InfinityStar shows a great advantage in terms of speed. Using visual tokenizers of the same compression rate, InfinityStar achieves a $10\times$ reduction in inference latency relative to leading diffusion models.

In summary, the main contributions of our work are as follows:

1. We propose InfinityStar, a novel spacetime pyramid modeling framework that unifies diverse visual generation tasks, demonstrating superior flexibility and versatility.
2. InfinityStar is the first discrete autoregressive model capable of generating high-quality videos, outperforming existing autoregressive text-to-video models and matching the performance of leading diffusion models.
3. Compared to the inefficiency of existing autoregressive models and diffusion models, InfinityStar significantly accelerates high-quality video generation.

2 Related Work

2.1 Video Diffusion Models

Diffusion models excel at generating high-fidelity data by gradually denoising random noise and has been widely applied in video generation. Early attempts [2, 4, 43] are built on U-Net architectures, demonstrating the feasibility of this approach but falling short in producing sharp, temporally coherent frames due to limited model capacity. The advent of Diffusion Transformers (DiT [23]) marked a turning point. SORA [3] harnessed DiT’s scaling ability to process spatio-temporal patches at scale, dramatically improving both video consistency and generation quality. The success of SORA has catalyzed a wave of innovation [40, 19, 32, 47] across the industry, propelling video generation to unprecedented levels of coherence and fidelity. Although video diffusion models deliver outstanding quality, their slow generation speed hinders the production of high-resolution, long-duration videos.

2.2 Video AutoRegressive Models

Another class of methods [34, 10, 18] employs autoregressive models for video generation. Inspired by the success of LLMs, these works predict video tokens in specific orders using an autoregressive Transformer. For example, Emu3 [34] performs next-token prediction along both spatial and temporal axes, while Nova [10] first predicts spatial tokens set-by-set and subsequently proceeds frame-by-frame in the temporal dimension. Although achieving preliminary progress, they require hundreds to thousands of inference steps, resulting in prohibitively low generation efficiency. In contrast, recent advances in next-scale prediction [29, 15] have demonstrated state-of-the-art performance in image synthesis, offering both improved quality and markedly faster inference. In this work, we extend the next-scale prediction paradigm to the unified image and video generation.

2.3 Discrete Video Tokenizers

For a long time, discrete [39, 41] and continuous [19, 40, 32] video tokenizers have been developed independently. Although some works [1, 33] provide both discrete and continuous tokenizers, the




Figure 1: **Spacetime pyramid modeling of InfinityStar.** Built with an unified autoregressive pipeline, InfinityStar is capable of performing text-to-image, text-to-video, image-to-video, video extrapolation tasks all in one model.

network configurations are usually not aligned. For example, Cosmos [1] chooses 6 and 16 as latent dimensions in its discrete and continuous variants respectively. This misalignment hinders the knowledge reuse between two types of tokenizers. As a result, most mainstream discrete video tokenizers are either trained from scratch [1] or starting from a pretrained discrete image tokenizer [41, 33]. However, these training strategies have the following drawbacks. First, training from scratch is inefficient and converges slowly. Second, weights pretrained on static images are not optimal for video reconstruction. To alleviate these deficiencies, we propose a new training strategy, which inherits the architecture and knowledge of a trained continuous video tokenizer. Experiments show that this strategy significantly boosts the convergence of discrete video tokenizers.

3 InfinityStar Architecture

3.1 Preliminaries

Infinity for Image Generation. Infinity [15] decomposes an image into a sequence of hierarchical token blocks using a visual tokenizer and models the relationship between tokens by a visual autoregressive Transformer (VAR Transformer). To cover images of various sizes, Infinity pre-defines a list of token block sizes $\{(h_1, w_1), \dots, (h_K, w_K)\}$, called scale schedule. The size (h_i, w_i) in scale schedule grows as i increases, forming a pyramid-like structure, which we refer as **image pyramid** in later discussion. Next we introduce the training and inference procedure of Infinity.

In the first training stage, a visual tokenizer learns to reconstruct the raw image and compress it into a sequence of discrete tokens, which can be modeled by the VAR Transformer in the next stage. Specifically, the tokenizer first encodes the raw images into compact latents, then transforms latents into K discrete residual token blocks $(r_1, r_2, \dots, r_K)$ using a bitwise multi-scale residual quantizer [15]. Each token block r_i consists of $h_i \times w_i$ discrete tokens of d -dim with vocabulary size of 2^d . Then in the second stage, a VAR Transformer is trained to predict next residual token block r_k conditioned on text embedding $\psi(t)$ and former tokens blocks $r_{<k}$. Formally, in each step, VAR Transformer predicts a conditional probability $p(r_k | r_{<k}, \psi(t))$. During the inference, Infinity generates an image by running the VAR Transformer K times autoregressively, merging the predicted tokens and running the tokenizer decoder once.




Figure 2: Influence of pretrained weights on reconstruction and convergence. The left sub-figure shows the reconstructed frames using different pretrained weights without `finetuning`. Loading weights of continuous video tokenizer achieves the best results. The right sub-figure shows that training with pretrained video tokenizer converges significantly faster than the other two strategies.

3.2 Spacetime Pyramid Modeling for Unified Generation

Extending the spatial-only next-scale prediction paradigm of Infinity [15] to video generation presents a primary challenge: *how to incorporate the temporal dimension*. The straightforward strategies are either letting time grows uniformly, *i.e.*, from $(1, 1, 1)$ to (T, H, W) , or keeping time constant, *i.e.*, from $(T, 1, 1)$ to (T, H, W) . We empirically found that letting time grow uniformly produces flickering videos. As for the constant time pyramid, we refer to it as the **pseudo-spacetime pyramid**. Despite its conceptual simplicity, it suffers from two fundamental limitations. First, the treatment of videos differs markedly from that of images, preventing a text-to-video (T2V) model from effectively leveraging the knowledge learned by a text-to-image (T2I) model and complicating its extension to tasks such as image-to-video (I2V). Second, because appearance and motion in videos are coupled in this design, the model faces significant difficulty in accurately fitting both aspects.

To overcome these challenges, we propose a novel **spacetime pyramid modeling** framework as shown in Fig. 1. Each video is decomposed into sequential clips $\{c_1, c_2, \dots, c_N\}$. We regard the first frame as c_1 (*i.e.*, $T = 1$) to specifically encode static appearance cues in videos and other clips share an equal duration $T > 1$. Each clip is modeled as a 3D volume pyramid similar to Infinity [15]. In particular, for each clip, there are K scales with each represented as a residual token block r_k of (T, h_k, w_k) dimension. *It is worth noting that all scales in the pyramid are extended only in spatial dimension instead of time*. Mathematically, the tokens in the first clip are generated auto-regressively across scales as:

$$p(r_1^1, \dots, r_K^1) = \prod_{k=1}^K p(r_k^1 | r_1^1, \dots, r_{k-1}^1, \psi(t)), \quad (1)$$

For inter-clip predictions, clips are generated sequentially conditioned on prior clip predictions and the text input in an autoregressive manner. In this way, we could generate infinitely long videos theoretically. Formally, the autoregressive likelihood of the whole video can be expressed as:

$$p(r_1^1, \dots, r_K^N) = \prod_{c=1}^N \prod_{k=1}^K p(r_k^c | r_1^1, \dots, r_{k-1}^c, \psi(t)), \quad (2)$$

3.3 Visual Tokenizer

Training video tokenizers faces greater challenges than training image tokenizers. First, training tokenizers on videos of tens of frames is much computationally heavier than training on static images. Therefore, training a video tokenizer from scratch is extremely time-consuming and suffers from slow convergence. Second, the scale schedule in videos leads to more imbalanced information distribution, where most information is concentrated in the last few scales. This brings great difficulties to the optimization of VAR Transformer. To solve these challenges, we introduce two techniques, knowledge inheritance from continuous video tokenizer and stochastic quantizer depth.




Figure 3: The influence of stochastic quantizer depth. Sub-figure (s_i, nt) represents the reconstructed frame nt using all tokens from the image pyramid plus tokens of first i scales in the clip pyramid. SQD significantly improves the reconstruction quality of early scales. Besides, the earlier scales correspond to global semantics, while the later ones are responsible for local visual details.

Knowledge Inheritance from Continuous Video Tokenizer. Instead of designing and training a discrete video tokenizer from scratch, we inherit the architecture and weights of a trained continuous video tokenizer, *i.e.* video VAE. Specifically, we first insert a parameter-free quantizer between the pre-trained VAE encoder and the decoder. The quantizer is based on binary spherical quantization [49], being similar to that of Infinity [15] but with new spacetime pyramid scale schedule. This does not introduce any new parameter like codebook in VQ [30] and well retains knowledge of the original VAE. As shown in Fig. 2, the discrete video tokenizer reconstructs videos decently, even without any fine-tuning. To further improve the reconstruction quality, we fine-tune the new tokenizer jointly on images and videos like previous works [33, 1]. During the fine-tuning, the KL loss of the original VAE is replaced with the commitment loss plus the entropy penalty [49]. As shown in Fig. 2, with the help of knowledge of continuous video VAE, the convergence accelerates dramatically.

Stochastic Quantizer Depth. When tokenizing videos using the spacetime pyramid schedule, the information distribution on different scales gets extremely imbalanced. Specifically, there are only a few tokens in the early scales, while there are tens of thousands of tokens in the last scales. Thus the tokenizer tends to reconstruct videos solely relying on tokens from the last few scales and not to learn useful representation in early scales as shown in Fig. 3 (left). However, this imbalanced distribution is difficult to model using VAR Transformer because the dependence between the latter token blocks and the former ones is weak. To alleviate this problem, we propose a regularization called stochastic quantizer depth. During training, each one of the last N scales has a probability p of being discarded. In this way, there are 2^N possible scale schedules during training. This requires the tokenizer to reduce the reliance on last scales and store more information in tokens of early scales. As in Fig. 3 (right), with the help of this regularization, the reconstruction results of early scales become much clearer. This balanced information distribution makes the training of VAR Transformer easier.

3.4 Spacetime Autoregressive Transformer

To accommodate the newly introduced temporal dimension, enhance the quality of generated videos, and alleviate the substantial computational overhead associated with a large number of tokens, we propose the following modifications to the VAR Transformer: Semantic Scale Repetition, Spacetime Sparse Attention, and Spacetime RoPE. We put Spacetime RoPE in the appendix A.

Semantic Scale Repetition. With carefully crafted positional encodings, InfinityStar can already generate videos of acceptable quality. However, we observe that the structural coherence and motion dynamics in these outputs remain suboptimal. As shown in Figure 3, the overall layout and the placement of foreground objects are determined by the early scales of the clip pyramid—what we term the “semantic scales.” This observation motivates us to enhance generation fidelity at these semantic scales. To this end, we introduce a simple yet effective technique called semantic scale repetition. Concretely, if a clip pyramid comprises K scale tuples, we repeat the first K_s tuples N times, thereby reinforcing the semantic representations. In this way, every early residual r_k undergoes multiple rounds of refinement, improving the generation quality of semantics and the performance in complex scenarios with large motion. Given that the tokens at these early scales account for only a small fraction of the total token count, the additional computational overhead incurred by repeating them is negligible.




Figure 4: Illustration of three causal attention variants. We plot three pyramids on the scale size = (1,2,3) for visualization simplicity. From left to right, VAR block-wise causal mask with full history, Switti block-wise non-causal mask with full history, and spacetime sparse attention.




Figure 5: Text to image and text to video examples.

Spacetime Sparse Attention. Autoregressive video generation faces significant challenges due to the high computational costs of long context. As on the left of Fig.4, Infinity [15] employs a block-wise causal mask for single pyramid modeling. Switti [31] verifies that conditioning solely on inputs from preceding scales is sufficient for next-scale predictions, resulting in a sparser attention mask as on the middle of Fig.4. For long video generation, it's necessary to attend history tokens to achieve temporal consistency. However, attending full history leads to an explosively long sequence. Considering each clip corresponds to 5s, which is sufficient to maintain temporal consistency, here we only attend to the last scale of the preceding clip. Finally, we obtain a highly sparse attention as show in Fig.4 (right). Our spacetime sparse attention drastically reduces computational overhead in attention during both training and inference, while delivering better performance.

4 Experiment

4.1 Implementation

Datasets. The training data of InfinityStar includes text-to-image data and text-to-video data. We curated 130M pretraining and 70M high-quality text-to-image data. To balance the data distribution and improve overall aesthetics, we also involve 5M high-quality synthetic data. In terms of text-to-video data, we curated around 16M video data. All videos are longer than 5 seconds. Among them 13M videos are under 336×192 resolution used for pre-training. They are mainly from Panda-70M[7], Mira[17], and other internal video-text pairs. Apart from those 192p videos, we also curated 3M 480p and 50K 720p high-quality videos for fine-tuning.

Model and Training. After inserting the patchify and unpatchify layers between Wan 2.1 VAE’s encoder and decoder, we obtain a video tokenizer with a compression rate of $4 \times 16 \times 16$ and a latent dimension of 64. Multi-scale BSQ quantization is adopted to obtain discrete tokens. In contrast to using a vocabulary size of 2^{64} for all scales, we use a vocabulary size of 2^{16} for the former small scales and 2^{64} for the latter large scales. We empirically find that it boosts convergence and has a negligible impact on the reconstruction quality. Starting with the pretrained weights of Wan 2.1 VAE, the discrete tokenizer is fine-tuned jointly on images of 256×256 , 512×512 , 768×768 and videos of $256 \times 256 \times 81$ for 30K iterations. The learning rate is $1e^{-4}$.

The autoregressive Transformer of InfinityStar is trained progressively in four stages, including a T2I pre-training and three T2V fine-tuning on 192p, 480p, 720p respectively. Each time we increase the training resolution, we preserve scale schedule of lower resolutions and append several larger scales, which enables better inheritance. The global batch size for 192p is 2048 and that of 480p and 720p is 1024. The learning rate for 192p is $2e^{-4}$. Then we decay it to $1e^{-4}$ for 480p and 720p. We train the model on videos of 192p, 480p, 720p for 50K, 8K, 3K iterations, respectively. Specifically, each clip pyramid is composed of 80 frames at 16 fps, and the first $K_s = 12$ semantic scales are repeated by $N = 3$ times. Details about infrastructure optimizations are presented in the appendix B.

4.2 Text-to-Image Generation

The upper part of Fig.5 shows images generated by our InfinityStar-T2I model, showcasing InfinityStar’s strength in generating high-fidelity and photo-realistic images across various categories and styles. We also carry out the quantitative evaluation on the GenEval[14] and DPG[16] benchmarks. As in Tab.1, InfinityStar achieves the best overall score of 0.79 on the GenEval bench with a prompt rewriter. It’s worth noting that InfinityStar exceeds Infinity by 6% on overall score. We attribute the significant improvement to the larger model size and the architectural innovations. On the DPG bench, InfinityStar reaches an overall score of 86.55, surpassing Infinity by 3.09%. These quantitative results demonstrate InfinityStar’s strong capabilities of image generation following users’ prompts.

4.3 Text-to-Video Generation

In the lower part of Fig.5, we present the generated videos of InfinityStar regarding user prompts. The generated videos successfully capture the semantic information in user prompts while maintaining high aesthetics and visual quality. Especially for the second example in Fig.5, the generated video accurately restores the delicate movements of the characters flipping through sketchbooks, talking while pointing to different parts of the drawings. In Tab.2, we compare InfinityStar with leading diffusion and autoregressive approaches on VBench—a comprehensive video benchmark spanning 16 evaluation dimensions. Our model achieves an overall score of 83.74, outperforming all open-source autoregressive baselines by a substantial margin. Moreover, InfinityStar surpasses diffusion-based competitors such as OpenSora[50], CogVideoX[40], and HunyuanVideo[19]. These results demonstrate that, through its novel spacetime autoregressive design, InfinityStar not only pushes the capabilities of discrete autoregressive video models but also attains performance on par with—and in some cases superior to—state-of-the-art diffusion methods.

Human Preference Evaluation. We conduct comprehensive human evaluation to benchmark our unified model, InfinityStar-8B, against a leading diffusion competitor, HunyuanVideo-13B. Specifically, we compared InfinityStar-8B to both the T2V and I2V variants of HunyuanVideo-13B. In a side-by-side comparison format, human raters were presented with videos generated by our model and those from HunyuanVideo-13B, and asked to judge which video was superior. Fig.6 lists the

Table 1: Evaluation on the GenEval [14] and DPG [16] benchmark. † result is with prompt rewriting or self-CoT.

Methods	# Params	GenEval↑				DPG↑		
		Two Obj.	Position	Color Attri.	Overall	Global	Relation	Overall
Diffusion Models								
SDXL [24]	2.6B	0.74	0.15	0.23	0.55	83.27	86.76	74.7
PixArt-Sigma [5]	0.6B	0.62	0.14	0.27	0.55	86.89	86.59	80.5
SD3 (d=38) [11]	8B	0.89	0.34	0.47	0.71	-	-	-
Goku [6]	2B	-	-	-	0.76†	-	-	83.6
Transfusion [51]	7.3B	-	-	-	0.63	-	-	-
SANA-1.0 [37]	1.6B	-	-	-	0.66	-	-	84.8
FLUX-dev [21]	12B	-	-	-	0.67	-	-	84.0
FLUX-schnell [21]	12B	-	-	-	0.71	-	-	84.8
AutoRegressive Models								
LlamaGen [26]	0.8B	0.34	0.07	0.04	0.32	-	-	65.2
Chameleon [27]	7B	-	-	-	0.39	-	-	-
Show-o [38]	1.3B	0.80	0.31	0.50	0.68	-	-	67.5
Liquid [36]	7B	0.73	0.17	0.37	0.55	-	-	-
UniTok [22]	7B	0.71	0.26	0.45	0.59	-	-	-
Janus [35]	1.3B	0.68	0.46	0.42	0.61	-	-	-
Emu3 [34]	8B	0.81†	0.49†	0.45†	0.66†	-	-	81.6
Fluid [12]	10.5B	0.83	0.39	0.51	0.69	-	-	-
NextStep-1 [28]	14B	-	-	-	0.73†	-	-	85.28
Infinity [15]	2B	0.85†	0.49†	0.57†	0.73†	93.11	90.76	83.46
InfinityStar-T2I	8B	0.90†	0.62†	0.67†	0.79†	91.68	91.87	86.55

Table 2: Evaluation on the VBench benchmark. † result is with prompt rewriting.

Models	# Params	Human Action	Scene	Multiple Objects	Appear. Style	Quality Score	Semantic Score	Overall
Diffusion Models								
AnimateDiff-V2	1.5B	92.60	50.19	36.88	22.42	82.90	69.75	80.27
VideoCrafter-2.0[4]	1.5B	95.00	<u>55.29</u>	40.66	25.13	82.20	73.42	80.44
OpenSora V1.2[50]	1.1B	85.80	42.47	58.41	23.89	80.71	73.30	79.23
Show-1[44]	6B	95.60	47.03	45.47	23.06	80.42	72.98	78.93
Gen-3 [13]	-	96.40	54.57	53.64	24.31	84.11	75.17	82.32
CogVideoX-5B[40]	5B	99.40	53.20	62.11	<u>24.91</u>	82.75	77.04	81.61
HunyuanVideo[19]	13B	94.40	53.88	68.55	19.80	85.09	75.82	83.24
Goku[6]	2B	97.60	57.08	<u>79.48</u>	23.08	<u>85.60</u>	81.87	84.85
Wan 2.1[32]	14B	<u>98.80</u>	53.67	81.44	21.13	85.64	<u>80.95</u>	<u>84.70</u>
AutoRegressive Models								
Nova[10]†	0.6B	95.20	54.06	77.52	20.92	80.39	79.05	80.12
Emu3[34]	8B	77.71	37.11	44.64	20.92	84.09	68.43	80.96
InfinityStar†	8B	96.43	52.08	<u>78.66</u>	21.81	<u>84.73</u>	<u>79.78</u>	<u>83.74</u>

results of two human preference benchmarks. For the T2V task, our model consistently outperformed HunyuanVideo-13B across all evaluation metrics, even while maintaining a notable speed advantage. For the I2V task, InfinityStar-8B also demonstrated superior performance, particularly in prompt following and overall quality, compared to HunyuanVideo-13B. These results highlight the robust capability of InfinityStar 8B in generating high-quality videos that adhere closely to textual prompts.

Zero-shot Generation. Although trained exclusively on T2V data, InfinityStar can generate videos conditioned on an image or a video as historical without any fine-tuning. Fig.7 shows video extrapolation results. The synthesized videos exhibit strong temporal coherence with the reference while faithfully capturing the semantic nuances of texts. Zero-shot I2V samples are presented in the appendix C.

4.4 Ablation Study

Visual Tokenizer. As shown in Fig.2 and Tab.3, loading weights of continuous video tokenizer significantly speeds up the convergence and achieves the best reconstruction results. As shown in Fig.3, stochastic quantizer depth largely improves the reconstruction quality of early scales. In terms of generation, using tokenizer with SQD leads to an improvement of 0.21 in VBench scores (81.28




Figure 6: Human evaluation results comparing our model with HunyuanVideo 13B.




Figure 7: Zero-shot video extrapolation examples. InfinityStar can extrapolate videos using a reference video as historical without any fine-tuning.

v.s. 81.07 as shown in Tab.4). Moreover, we observe that SQD contributes to faster convergence during the video generation training.

Pseudo-Spacetime Pyramid v.s. Spacetime Pyramid.

As illustrated in Fig.8, videos generated by the pseudo-spacetime pyramid lack visual details and deliver simpler motion. In contrast, spacetime pyramid generates videos with richer details and higher motion. Besides, spacetime pyramid improves VBench's overall score from 80.30 to 81.28 as illustrated in Tab.4. These experiments support the hypothesis that spacetime pyramid could decouple

Table 3: Reconstruction metrics on an internal high-motion video benchmark (480p 81 frames).

Pretrained Weights	PSNR($\uparrow$)	SSIM($\uparrow$)	LPIPS($\downarrow$)
Continuous Video VAE	33.37	0.94	0.065
Image VAE	29.10	0.90	0.123
None	30.04	0.90	0.124



Figure 8: Comparison between Pseudo-Spacetime Pyramid and Spacetime Pyramid. Spacetime Pyramid could generate videos with richer details and higher motion.

appearance and temporal information. The image pyramid corresponds to the appearance information and clip pyramids focus on subsequent motions. This decoupling makes it easier to learn motions in videos. In addition to advances in performance, spacetime pyramid unifies T2I, T2V, I2V tasks into one framework.

Semantic Scale Repetition. In Fig.3, we can observe that the earlier scales correspond to semantic information, while the latter ones are responsible for high-frequency details. Here we compare the generation results of with and without semantic scale repetition. As shown in Fig.9, semantic scale repetition is highly effective in improving the structure stability and motion quality. The quantitative results further confirm the significant gains. As shown in Tab.4, semantic scale repetition improves VBench’s overall score from 75.72 to 81.28.

Spacetime Sparse Attention. In Tab.4 and Tab.5, we compare different attention mechanisms. Spacetime sparse attention shows superior performance to full attention in the VBench total score (81.28 v.s. 80.77), while showing a significant advantage in saving computation and GPU VRAM. SSA reaches $1.5\times$ speedup when generating 192p 161 frames. The efficiency advantage becomes larger as the resolution and duration grow. For 480p 161 frames, full attention fails due to OOM while SSA completes it within 44.7s using 63GB VRAM. We hypothesize that SSA produces better results than full attention because it reduces exposure bias. Full attention is more susceptible to accumulated errors. The reason we do not condition on smaller scales of the preceding clip is that it misses the former clips’ visual details and brings visual inconsistency between clips. Although it reaches $1.1\times$, $1.5\times$ speedup for 192p and 480p 161 frames, we observe a significant performance drop in VBench from 81.28 to 80.75 as shown in Tab.4. Therefore, the proposed spacetime sparse attention strikes a better balance between computational efficiency and visual quality.

4.5 Inferency Latency

As shown in Tab.6, we report the end-to-end inference latency measured on a single GPU, including both the text encoder and VAE decoder. Wan-2.1[32] and Nova[10] were evaluated using their default GitHub configurations. Even without employing stronger compression, InfinityStar achieves a $32\times$




Figure 9: Semantic Scale Repetition (SSR) greatly improves structure stability and motion quality.

speedup over Wan-2.1. Furthermore, despite its $13\times$ larger model size, InfinityStar delivers a $6\times$ speedup compared to Nova. These results highlight our model’s significant advantage in efficiency over both diffusion and autoregressive approaches.

Table 4: Comprehensive ablation studies. Experiment with 1M 192p training data, batch size of 40, and 30K iterations. We evaluate the results on the Vbench benchmark.

Vbench	total score	quality score	semantic score
InfinityStar (Our Model)	81.28	81.56	80.16
<i>Attend to former clip’s largest scale</i>			
<i>Ablation by removing/replacing core components</i>			
w/o Semantic Scale Repetition(SSR)	75.72	76.73	71.68
w/o Spacetime Pyramid (using Pseudo-Spacetime)	80.30	80.81	78.28
w/o Stochastic Quantizer Depth(SQD)	81.07	81.21	80.54
<i>Comparison of different Attention Mechanism variants</i>			
Full Attention	80.77	81.15	79.23
Attend to former clip’s 3rd largest scale	80.86	81.26	79.26
Attend to former clip’s 6th largest scale	80.75	80.98	79.80

Table 5: Computational efficiency comparison of attention mechanisms on a single GPU.

	(192p 65 frames)	(192p 161 frames)	(480p 161 frames)
Full Attention	8.6s / 40.8GB	24.3s / 57GB	OOM
Attend to former clip’s largest scale	7.7s / 38.5GB	16.7s / 40GB	44.7s / 63 GB
Attend to former clip’s 3rd largest scale	7.4s / 38.2GB	15.8s / 39GB	34.5s / 58 GB
Attend to former clip’s 6th largest scale	7.3s / 37.9GB	15.2s / 38GB	30.5s / 55GB

Table 6: Computational efficiency comparison.

Method	Model	# Parameters	Durations(s)	Frames	Resolution	Time(s)	Speedup
Diffusion	Wan 2.1[32]	14B	5	81	720p	1864	1
AR	Nova[10]	0.6B	5	81	480p	354	5
AR	InfinityStar	8B	5	81	720p	58	32




Figure 10: **Framework of InfinityStar-Interact.** We propose Semantic-Detail conditions (illustrated in light blue cubes) to control video synthesizing when interacting with users. It delivers superior visual and semantic consistency, as well as strong prompt-following capabilities.

5 Extended Application: Long Interactive Video Generation

The long interactive video generation task focuses on the collaborative generation between the T2V model and users, accepting new user instructions and generating corresponding video content iteratively. While the original InfinityStar supports generating 10-second 480p videos, it only accepts a single prompt input and is limited to two clips. Extrapolating to longer video lengths than training involves performance degradation due to the discrepancy between training and testing. Simply increasing the number of training clips will lead to excessively long training sequences, which in turn causes an OOM issue. Below we introduce the innovations to extend InfinityStar to support long interactive video generation.

5.1 Model Design

We solve the problem of long interactive video generation using a sliding window method. Mathematically, for a long interactive video $V \in T^{long} \times H \times W$, we decompose it into a series of video chunks of 10 seconds, *i.e.*, $\{V_0, V_1, \dots, V_n\}$, with stride of 5 seconds. Each chunk V_i is further divided into two clips V_i^0 and V_i^1 . Each video clip is 5 seconds long and paired with a transition caption, *i.e.*, t_{i-1}^1 or t_i^0 , with the assistance of an LLM. Note that (t_i^0, V_i^0) is the same with (t_{i-1}^1, V_{i-1}^1) . During each round interaction with the user, InfinityStar generates V_i^1 conditioned on $(V_0^0[0, \dots], V_i^0, t_i^1)$, where V_i^0 is V_{i-1}^1 that we generated in the preceding interaction round. $V_0^0[0, \dots]$ is the first frame of the earliest video clip. This division method allows training on only two clips, while enabling to synthesize infinitely long videos during the inference stage. We find that conditioning on $V_0^0[0, \dots]$ could mitigate drift when generating multi-round videos.

Beyond spacetime sparse attention, we introduce the novel Semantic-Detail conditions to control video synthesizing when interacting with users as illustrated in Fig.10. Specifically, we extract features $F_{i-1} \in T \times H \times W$ from the preceding clip V_{i-1}^1 using the visual tokenizer. The features F_{i-1} are referred to detail features since they are full-scale and contain rich visual details. It is difficult to extract semantic information from F_{i-1} because it is not adequately compressed. Besides, there are too many tokens in F_{i-1} , which significantly slows down the interactive inference speed. Borrow ideas from FramePack [46], we downsample F_{i-1} to $F_{i-1}^{sem} \in T \times h \times w$ spatially to reduce excessive condition tokens. The semantic conditions F_{i-1}^{sem} are employed to enable semantic consistency between clips. Apart from F_{i-1}^{sem} , we slice the last K frames from F_{i-1} instead of the whole as the detail conditions $F_{i-1}^{det} \in K \times H \times W$. In this way, we ensure consistency in both semantics and details while significantly compressing the number of condition tokens.




Figure 11: Conditioning solely on the last few frames of preceding clip (baseline conditions) is inadequate for preserving semantic consistency. Our proposed conditions deliver better capability in maintaining semantic consistency.

5.2 Dataset

We curate the interactive generation data from the pretraining dataset and other sources. In particular, we select videos longer than 7 seconds from the pretraining data, resulting in a total of 7M videos. Subsequently, we decompose long videos into chunks, split the chunks into clips, and generate captions at the clip level using the Tarsier2 [42] model. It is worth noting that here we adopt an LLM to remove the content that had already appeared in V_i^0 's caption from V_i^1 's caption, and ensure that t_i^1 only describes changes compared to t_i^0 . The instructions used to query an LLM are presented in the Appendix C.4. In this way, we align with the instructions users provide during interactive generation.

Apart from filtering pretraining data, we also incorporate some synthetic long interaction data. Specifically, we first collect multi-round interactive prompts. These prompts are used as seeds to query an LLM to generate more samples. We pick good ones from the generated samples to enlarge the seed set and query an LLM again to enhance diversity. Finally, we collect 16K interactive prompts, where each prompt consists of four round interactions. Then we use the prompts to query a video continuation model to generate interaction videos. We provide the instructions to generate multi-round interactive prompts in the Appendix C.4. We present some examples of the curated interaction data in Fig. 12.

5.3 Evaluation

The training of the interactive generation model is divided into two stages. In the first stage, we load the weights of InfinityStar and conduct continued pre-training on the filtered pre-training data. The learning rate during this stage is set to $2e-4$. In the second stage, we fine-tune the model on the synthetic interaction data. We decay the learning rate to $2e-5$. We slice the last 2 frames (set $K = 2$) from the preceding clip as detail features. The semantic features are obtained by downsampling the detail features with a stride of $\sqrt{32}$. Compared to spacetime sparse attention, the proposed semantic-detail conditions compress the condition token length from 33.6K to 5.8K for 480P video generation.

Empirical observations reveal that relying solely on the last few frames of the preceding clip (abbreviated as baseline conditions) is inadequate to preserve semantic consistency in the long interactive generation task. Our proposed semantic-detail conditions deliver higher quality and better consistency in semantics while showing high efficiency. As shown in Fig. 11, the face ID of the woman has changed after three rounds of interactive generation, whereas the proposed conditions have successfully maintained its consistency. Fig. 13 presents two examples of InfinityStar-Interact. Whether outdoor character movements as in the first example or indoor character hand movements as in the second example, InfinityStar-Interact generates consistent videos during interactions with the user.




Figure 12: Examples of curated interactive training data. The upper part is obtained by selecting data from pre-training datasets and rewriting captions using an LLM. The lower part is synthetic interaction data, generated by first using an LLM to create prompts and then calling a video continuation model.

6 Conclusion

We introduce InfinityStar, a unified spacetime autoregressive framework capable of synthesizing high-resolution images and dynamic, high-motion videos. By seamlessly integrating spatial and temporal prediction within a purely discrete architecture, InfinityStar supports diverse generation tasks while maintaining both state-of-the-art quality and exceptional efficiency. Our extensive evaluation demonstrates that InfinityStar outperforms prior autoregressive video models and rivals leading diffusion-based approaches, producing a 720p video of 5s in one-tenth the inference time. Besides, we extend InfinityStar to support long interactive video generation. As the first discrete autoregressive model to deliver industrial-grade 720p video synthesis, we anticipate that InfinityStar will catalyze future research on rapid, long video generation.

7 Limitation

While InfinityStar sets a new record in discrete video generation and demonstrates strong prompt following ability as well as impressive motion capabilities, several limitations remain. Specifically, there is a trade-off between image quality and motion fidelity in high-motion scenes, where sometimes fine-grained visual details can be compromised. Additionally, due to limited computational resources, we have not scaled our model training or parameter size to match those of leading diffusion models, which constrains the upper bound of the performance. Furthermore, our inference pipeline has not yet been fully optimized, indicating room for future improvement. In terms of the limitations in long interactive video generation, InfinityStar suffers from cumulative errors. With the increase in the number of interactions, there will be a noticeable degradation in the quality of the generated videos. This constitutes a problem that we are required to address.

8 Acknowledgments

The authors appreciate the valuable support provided by colleagues from ByteDance, including Yuqi Zhang, Yifu Zhang, Hao Yang, Yifei Hu, Chuang Lin, Xiaofeng Mei, Ruibiao Lu, and Jiawei Duan. Their contributions to data processing and related technical aspects are essential for the advancement of this research.

[0s~5s] Reference



[5s~10s] The boy bends down and picks up the red soccer ball from the bench.



[10s~15s] The boy holds the soccer ball with both hands and begins to bounce it on the ground.



[15s~20s] The boy kicks the soccer ball forward, sending it rolling across the grass.



[0s~5s] Reference



[5s~10s] The man begins slicing a red bell pepper on the cutting board with the knife in his right hand.



[10s~15s] The man picks up a wooden spoon from the countertop and stirs the contents of a pot on the stove.



[15s~20s] The man wipes his hands on a kitchen towel hanging from the oven handle while glancing at the pot.



Figure 13: Interactive Generation Results. Given the first 5-second video as a reference, InfinityStar-Interact generates 480p videos through multi-round collaboration with users. Whether focusing on outdoor character movements (as in the first example) or indoor character hand movements (as in the second example), InfinityStar-Interact can generate interactive videos that follow users' prompts.

References

- [1] N. Agarwal, A. Ali, M. Bala, Y. Balaji, E. Barker, T. Cai, P. Chattopadhyay, Y. Chen, Y. Cui, Y. Ding, et al. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025. [2](#), [3](#), [5](#)
- [2] A. Blattmann, T. Dockhorn, S. Kulal, D. Mendelevitch, M. Kilian, D. Lorenz, Y. Levi, Z. English, V. Voleti, A. Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. [2](#)
- [3] T. Brooks, B. Peebles, C. Holmes, W. DePue, Y. Guo, L. Jing, D. Schnurr, J. Taylor, T. Luhman, E. Luhman, C. Ng, R. Wang, and A. Ramesh. Video generation models as world simulators. *OpenAI*, 2024. [1](#), [2](#)
- [4] H. Chen, M. Xia, Y. He, Y. Zhang, X. Cun, S. Yang, J. Xing, Y. Liu, Q. Chen, X. Wang, et al. Videocrafter1: Open diffusion models for high-quality video generation. *arXiv preprint arXiv:2310.19512*, 2023. [2](#), [8](#)
- [5] J. Chen, C. Ge, E. Xie, Y. Wu, L. Yao, X. Ren, Z. Wang, P. Luo, H. Lu, and Z. Li. Pixart-sigma: Weak-to-strong training of diffusion transformer for 4k text-to-image generation. *arXiv preprint arXiv:2403.04692*, 2024. [8](#)
- [6] S. Chen, C. Ge, Y. Zhang, Y. Zhang, F. Zhu, H. Yang, H. Hao, H. Wu, Z. Lai, Y. Hu, et al. Goku: Flow based video generative foundation models. *arXiv preprint arXiv:2502.04896*, 2025. [8](#), [21](#)
- [7] T.-S. Chen, A. Siarohin, W. Menapace, E. Deyneka, H.-w. Chao, B. E. Jeon, Y. Fang, H.-Y. Lee, J. Ren, M.-H. Yang, et al. Panda-70m: Captioning 70m videos with multiple cross-modality teachers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13320–13331, 2024. [7](#)
- [8] Z. Chen, J. Wu, W. Wang, W. Su, G. Chen, S. Xing, Z. Muyan, Q. Zhang, X. Zhu, L. Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv preprint arXiv:2312.14238*, 2023. [20](#)
- [9] G. DeepMind. Veo 3. <https://deepmind.google/technologies/veo/veo-3/>, 2025.05. [1](#)
- [10] H. Deng, T. Pan, H. Diao, Z. Luo, Y. Cui, H. Lu, S. Shan, Y. Qi, and X. Wang. Autoregressive video generation without vector quantization. *arXiv preprint arXiv:2412.14169*, 2024. [1](#), [2](#), [8](#), [10](#), [11](#)
- [11] P. Esser, S. Kulal, A. Blattmann, R. Entezari, J. Müller, H. Saini, Y. Levi, D. Lorenz, A. Sauer, F. Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. [8](#)
- [12] L. Fan, T. Li, S. Qin, Y. Li, C. Sun, M. Rubinstein, D. Sun, K. He, and Y. Tian. Fluid: Scaling autoregressive text-to-image generative models with continuous tokens. *arXiv preprint arXiv:2410.13863*, 2024. [8](#)
- [13] A. Germanidis. Introducing gen-3 alpha: A new frontier for video generation. *Runway*, 2024. [8](#)
- [14] D. Ghosh, H. Hajishirzi, and L. Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems*, 36, 2024. [7](#), [8](#)
- [15] J. Han, J. Liu, Y. Jiang, B. Yan, Y. Zhang, Z. Yuan, B. Peng, and X. Liu. Infinity: Scaling bitwise autoregressive modeling for high-resolution image synthesis. *arXiv preprint arXiv:2412.04431*, 2024. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#), [19](#)
- [16] X. Hu, R. Wang, Y. Fang, B. Fu, P. Cheng, and G. Yu. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024. [7](#), [8](#)
- [17] X. Ju, Y. Gao, Z. Zhang, Z. Yuan, X. Wang, A. Zeng, Y. Xiong, Q. Xu, and Y. Shan. Miradata: A large-scale video dataset with long durations and structured captions, 2024. [7](#)
- [18] D. Kondratyuk, L. Yu, X. Gu, J. Lezama, J. Huang, G. Schindler, R. Hornung, V. Birodkar, J. Yan, M.-C. Chiu, K. Somandepalli, H. Akbari, Y. Alon, Y. Cheng, J. Dillon, A. Gupta, M. Hahn, A. Hauth, D. Hendon, A. Martinez, D. Minnen, M. Sirotenko, K. Sohn, X. Yang, H. Adam, M.-H. Yang, I. Essa, H. Wang, D. A. Ross, B. Seybold, and L. Jiang. Videopoet: A large language model for zero-shot video generation, 2024. [1](#), [2](#)
- [19] W. Kong, Q. Tian, Z. Zhang, R. Min, Z. Dai, J. Zhou, J. Xiong, X. Li, B. Wu, J. Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024. [1](#), [2](#), [7](#), [8](#)
- [20] Kuaishou. Kling ai. <https://klingai.kuaishou.com/>, 2024.06. [1](#)
- [21] B. F. Labs. Flux. <https://blackforestlabs.ai/announcing-black-forest-labs/>, 2024. [8](#)
- [22] C. Ma, Y. Jiang, J. Wu, J. Yang, X. Yu, Z. Yuan, B. Peng, and X. Qi. Unitok: A unified tokenizer for visual generation and understanding. *arXiv preprint arXiv:2502.20321*, 2025. [8](#)
- [23] W. Peebles and S. Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. [2](#)
- [24] D. Podell, Z. English, K. Lacey, A. Blattmann, T. Dockhorn, J. Müller, J. Penna, and R. Rombach. Sd-xl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. [8](#)
- [25] J. Su, M. Ahmed, Y. Lu, S. Pan, W. Bo, and Y. Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. [19](#)
- [26] P. Sun, Y. Jiang, S. Chen, S. Zhang, B. Peng, P. Luo, and Z. Yuan. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*, 2024. [8](#)
- [27] C. Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024. [8](#)
- [28] N. Team, C. Han, G. Li, J. Wu, Q. Sun, Y. Cai, Y. Peng, Z. Ge, D. Zhou, H. Tang, H. Zhou, K. Liu, A. Huang, B. Wang, C. Miao, D. Sun, E. Yu, F. Yin, G. Yu, H. Nie, H. Lv, H. Hu, J. Wang, J. Zhou, J. Sun, K. Tan, K. An, K. Lin, L. Zhao, M. Chen, P. Xing, R. Wang, S. Liu, S. Xia, T. You, W. Ji, X. Zeng, X. Han, X. Zhang, Y. Wei, Y. Xu, Y. Jiang, Y. Wang, Y. Zhou, Y. Han, Z. Meng, B. Jiao, D. Jiang, X. Zhang, and Y. Zhu. Nextstep-1: Toward autoregressive image generation with continuous tokens at scale. *arXiv*

- preprint *arXiv:2508.10711*, 2025. 8
- [29] K. Tian, Y. Jiang, Z. Yuan, B. Peng, and L. Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *arXiv preprint arXiv:2404.02905*, 2024. 1, 2
 - [30] A. Van Den Oord, O. Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 5
 - [31] A. Voronov, D. Kuznedelev, M. Khoroshikh, V. Khrulkov, and D. Baranchuk. Switti: Designing scale-wise transformers for text-to-image synthesis. *arXiv preprint arXiv:2412.01819*, 2024. 6
 - [32] A. Wang, B. Ai, B. Wen, C. Mao, C.-W. Xie, D. Chen, F. Yu, H. Zhao, J. Yang, J. Zeng, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 1, 2, 8, 10, 11
 - [33] J. Wang, Y. Jiang, Z. Yuan, B. Peng, Z. Wu, and Y.-G. Jiang. Omnitokenizer: A joint image-video tokenizer for visual generation. *Advances in Neural Information Processing Systems*, 37:28281–28295, 2024. 2, 3, 5
 - [34] X. Wang, X. Zhang, Z. Luo, Q. Sun, Y. Cui, J. Wang, F. Zhang, Y. Wang, Z. Li, Q. Yu, et al. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*, 2024. 1, 2, 8
 - [35] C. Wu, X. Chen, Z. Wu, Y. Ma, X. Liu, Z. Pan, W. Liu, Z. Xie, X. Yu, C. Ruan, et al. Janus: Decoupling visual encoding for unified multimodal understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 12966–12977, 2025. 8
 - [36] J. Wu, Y. Jiang, C. Ma, Y. Liu, H. Zhao, Z. Yuan, S. Bai, and X. Bai. Liquid: Language models are scalable and unified multi-modal generators. *arXiv preprint arXiv:2412.04332*, 2024. 8
 - [37] E. Xie, J. Chen, J. Chen, H. Cai, H. Tang, Y. Lin, Z. Zhang, M. Li, L. Zhu, Y. Lu, et al. Sana: Efficient high-resolution image synthesis with linear diffusion transformers. *arXiv preprint arXiv:2410.10629*, 2024. 8
 - [38] J. Xie, W. Mao, Z. Bai, D. J. Zhang, W. Wang, K. Q. Lin, Y. Gu, Z. Chen, Z. Yang, and M. Z. Shou. Show-o: One single transformer to unify multimodal understanding and generation. *arXiv preprint arXiv:2408.12528*, 2024. 8
 - [39] W. Yan, Y. Zhang, P. Abbeel, and A. Srinivas. Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021. 2
 - [40] Z. Yang, J. Teng, W. Zheng, M. Ding, S. Huang, J. Xu, Y. Yang, W. Hong, X. Zhang, G. Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 2, 7, 8
 - [41] L. Yu, J. Lezama, N. B. Gundavarapu, L. Versari, K. Sohn, D. Minnen, Y. Cheng, A. Gupta, X. Gu, A. G. Hauptmann, et al. Language model beats diffusion—tokenizer is key to visual generation. *arXiv preprint arXiv:2310.05737*, 2023. 2, 3
 - [42] L. Yuan, J. Wang, H. Sun, Y. Zhang, and Y. Lin. Tarsier2: Advancing large vision-language models from detailed video description to comprehensive video understanding. *arXiv preprint arXiv:2501.07888*, 2025. 13, 20
 - [43] Y. Zeng, G. Wei, J. Zheng, J. Zou, Y. Wei, Y. Zhang, and H. Li. Make pixels dance: High-dynamic video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8850–8860, 2024. 2
 - [44] D. J. Zhang, J. Z. Wu, J.-W. Liu, R. Zhao, L. Ran, Y. Gu, D. Gao, and M. Z. Shou. Show-1: Marrying pixel and latent diffusion models for text-to-video generation. *International Journal of Computer Vision*, pages 1–15, 2024. 8
 - [45] F. Zhang, S. Tian, Z. Huang, Y. Qiao, and Z. Liu. Evaluation agent: Efficient and promptable evaluation framework for visual generative models. *arXiv preprint arXiv:2412.09645*, 2024. 2
 - [46] L. Zhang and M. Agrawala. Packing input frame context in next-frame prediction models for video generation. *arXiv preprint arXiv:2504.12626*, 2025. 12
 - [47] Y. Zhang, H. Yang, Y. Zhang, Y. Hu, F. Zhu, C. Lin, X. Mei, Y. Jiang, Z. Yuan, and B. Peng. Waver: Wave your way to lifelike video generation. *arXiv preprint arXiv:2508.15761*, 2025. 1, 2
 - [48] Y. Zhao, A. Gu, R. Varma, L. Luo, C.-C. Huang, M. Xu, L. Wright, H. Shojanazeri, M. Ott, S. Shleifer, et al. Pytorch fsdp: experiences on scaling fully sharded data parallel. *arXiv preprint arXiv:2304.11277*, 2023. 20
 - [49] Y. Zhao, Y. Xiong, and P. Krähenbühl. Image and video tokenization with binary spherical quantization. *arXiv preprint arXiv:2406.07548*, 2024. 5
 - [50] Z. Zheng, X. Peng, T. Yang, C. Shen, S. Li, H. Liu, Y. Zhou, T. Li, and Y. You. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*, 2024. 7, 8
 - [51] C. Zhou, L. Yu, A. Babu, K. Tirumala, M. Yasunaga, L. Shamsi, J. Kahn, X. Ma, L. Zettlemoyer, and O. Levy. Transfusion: Predict the next token and diffuse images with one multi-modal model. *arXiv preprint arXiv:2408.11039*, 2024. 8

A Spacetime Autogressive Modeling

Spacetime RoPE. We introduce spacetime rotary position embeddings (Spacetime RoPE) tailored for InfinityStar. This is achieved by decomposing original rotary embeddings[25] into four components: scale, time, height, and width. As shown in Fig. 14, the scale ID serves as a counter of scales up to now. The temporal ID remains zero for tokens in the image pyramid. For tokens in video pyramids, it increments as the frame grows. Distinct IDs are assigned to height and width components based on the token’s position in the image or video. Spacetime RoPE enhances the modeling of complex positional information for tokens in image and video pyramids.

Spacetime Autoregressive Transformer with Bitwise Self-Correction. To alleviate the train-test discrepancies of teacher-forcing training, we adopt bitwise self-correction mechanism proposed by Infinity[15]. Specifically, during training, some of the input tokens are randomly flipped to simulate the prediction error during the inference phase. Besides, the target labels are also recomputed to match the perturbed inputs. Moreover, when predicting the token distribution, the traditional index-wise classifier is replaced by a bitwise classifier. The bitwise classifier predicts d bits instead of 2^d indices, significantly reducing the memory costs and difficulties in optimization. Algorithm 1 shows the detailed procedure of Spacetime Pyramid Encoding with Bitwise Self-Correction.

Algorithm 1 Spacetime Pyramid Encoding with BSC

Input: raw feature F , scale schedule number K , clip number N
 image pyramid scale schedule: $(1, h_1, w_1), \dots, (1, h_K, w_K)$,
 clip pyramid scale schedule: $(T, h_1, w_1), \dots, (T, h_K, w_K)$

$R_{queue} \leftarrow []$ $\triangleright$ multi-scale bit labels
 $\tilde{F}_{queue} \leftarrow []$ $\triangleright$ inputs for transformer
for $c = 1, 2, \dots, N$ **do** $\triangleright$ inter-clips iteration
 $t_{start} = 1 + (c - 1) * T$
 $F_c \leftarrow$ raw features from time t_{start} to $t_{start} + t_c$
 for $k = 1, 2, \dots, K$ **do** $\triangleright$ intra-clip multi-scale iteration
 $R_k = \text{quant}(\text{down}(F_c - F_{c,k-1}^{flip}, (t_k, h_k, w_k)))$
 Queue_Push(R_{queue}, R_k)
 $R_k^{flip} = \text{Random_Flip}(R_k, p)$
 $F_{c,k}^{flip} = \sum_{i=1}^k \text{up}(R_i^{flip}, (h, w))$
 $\tilde{F}_{c,k} = \text{down}(F_{c,k}^{flip}, (t_{k+1}, h_{k+1}, w_{k+1}))$
 Queue_Push($\tilde{F}_{queue}, \tilde{F}_{c,k}$)
end for
end for
Output: $R_{queue}, \tilde{F}_{queue}$




Figure 14: An illustration of Spacetime RoPE. We decompose rotary embeddings into four components, i.e., scale, time, height, and width components. Spacetime RoPE enhances the modeling of complex positional information while supporting extrapolation.

B Infrastructure and Data

Infrastructure Optimization. Compared to diffusion models, visual autoregressive methods possess around $2.5\times$ longer training sequences. This feature poses crucial pressure on hardware and algorithms when scaling models and increasing resolutions. In this work, we adopt advanced parallelism methods for scalable and efficient training.




Figure 15: Text to image examples.

Firstly, we utilize FlexAttention to implement various attention mechanisms. With our proposed Spacetime Sparse Attention, we achieve more than a $2\times$ acceleration in training speed. Secondly, we adopt fully sharded data parallelism (FSDP) [48] to partition parameters, gradients, and optimizer states across GPU ranks. Thirdly, we adopt a fine-grained activation checkpointing strategy to reduce the overhead of vRAM and data transfer, making the parallelization more efficient. Last but not least, sequence parallelism further partitions long sequences into multiple chunks and then exploits ring self-attention for each chunk, making it feasible to train 720p videos with 200K sequence length.

Visual Captioning. Detailed visual captioning is crucial for enabling the model to accurately generate images and videos. For images, we use InternVL2.0[8] to produce dense descriptions for each sample. For video clips, we obtain overall video descriptions using Tarsier2[42]. Notably, Tarsier2 inherently captures camera motion types (e.g., zoom, pan right), eliminating the need for a separate prediction model. This simplifies the pipeline and improves efficiency.




Figure 16: Zero-shot image to video examples. InfinityStar can generate videos following an input image without fine-tuning. The synthesized videos exhibit strong temporal and semantic coherence.

Data Pipeline. Obtaining a high-quality image and video dataset requires a complex processing pipeline. Specifically for video, we follow video processing pipelines[6] to preprocess videos into high-quality training clips through OCR filtering, video clip extraction, visual aesthetic filtering, and motion filtering, etc.

C More Qualitative Results

C.1 Text-to-Image Generation.

Fig.15 shows more generated images from our InfinityStar-T2I model. Our model is capable of generating high-resolution images filled with vivid and intricate details.

C.2 Zero-shot Generation

Image to Video. Although trained exclusively on text-to-video data, InfinityStar can generate videos conditioned on an input image without any fine-tuning. Fig.16 illustrates qualitative results on the image-to-video task. The synthesized videos exhibit strong temporal coherence with the reference image—a critical requirement for this task—while faithfully capturing the semantic nuances of the accompanying text with high visual fidelity.

C.3 Video Reconstruction.

Figure 17 illustrates a comparison between the reconstructed videos generated by different tokenizers and the original video. The discrete tokenizer trained from scratch (middle row) exhibits inferior reconstruction quality. In contrast, the tokenizer incorporating knowledge inheritance (top row) demonstrates a substantial improvement in visual fidelity, particularly in the preservation of intricate details such as human faces and complex textures.



Figure 17: Comparison between the reconstruction quality of different video tokenizers. The tokenizer incorporating knowledge inheritance (top row) demonstrates a substantial improvement compared to one trained from scratch (middle row).

C.4 Instructions.

Below is the instruction for removing duplicate captions from adjacent clips.

You are a helpful assistant.

Paragraph 1: <<<clip 1's tarsier2 caption>>>

Paragraph 2: <<<clip 2's tarsier2 caption>>>

These two paragraphs describe a 10-second video: the first paragraph covers the first 5 seconds, while the second focuses on the last 5 seconds.

However, the second paragraph was written without considering the content already included in the first one, resulting in significant repetition.

Now, I need you to revise the second paragraph:

- Remove the repetitive content that has already been mentioned in the first paragraph and retain only the new information.
- You can think of the revised second paragraph as a description of what changes occurred in the last 5 seconds compared to the first 5 seconds.
- If necessary, add sequential transition words such as "then" or "next" to better describe the changes.
- If no obvious differences are identified, you may first extract the core content from the previous paragraph and then add transition words like "continue" or "keep" to indicate continuity.
- Please provide an analysis first, followed by the revised result.
- Please place the revised results between "<<<" and ">>>"

Below is the instruction for generating multi-round interactive prompts.

You are an expert in writing prompts. The written prompts are used to query a text-to-video model to generate videos interactively. Each video is 20 seconds long and consists of four 5-second shots. Each shot shows the next moment of the same scene compared to the previous shot. For each new shot, you add a new action to the main subject from the previous shot. Describe the facts directly and do not use rhetoric. To prevent hallucinations, the objects in the subsequent three shots must have appeared in the first shot.

Below are some examples you have written before:

Example 1

<story>

<shot1>A young boy wearing a green hoodie and jeans is in a backyard with a wooden fence and green grass. A red ball, a blue bicycle, and a yellow toy truck are on the grass nearby. The boy is standing next to the red ball, looking at it with his hands on his hips.</shot1>

<shot2>The boy picks up the red ball with both hands.</shot2>

<shot3>The boy throws the red ball forward across the grass.</shot3>

<shot4>The boy runs toward the blue bicycle parked near the fence.</shot4>

</story>

Example 2

<story>
<shot1>A woman wearing a red sweater and glasses stands in a kitchen with white cabinets and a marble countertop. On the countertop are a cutting board with chopped vegetables, a stainless steel knife, a glass bowl, and a bottle of olive oil. The woman holds the knife in her right hand and is about to chop a tomato on the cutting board.</shot1>
<shot2>The woman finishes chopping the tomato and places the knife down on the cutting board.</shot2>
<shot3>The woman picks up the glass bowl and transfers the chopped vegetables into it.</shot3>
<shot4>The woman picks up the bottle of olive oil and pours some into the glass bowl.</shot4>
</story>

Example 3

<story>
<shot1>A man wearing a blue button-up shirt and black trousers stands in a small home office. The room contains a wooden bookshelf filled with books, a black swivel chair, and a desk with a desktop computer, a white coffee mug, and a closed notebook. The man holds a smartphone in his right hand, looking at the screen with a neutral expression.</shot1>
<shot2>The man puts the smartphone down on the desk next to the coffee mug.</shot2>
<shot3>The man sits down on the black swivel chair and opens the notebook on the desk.</shot3>
<shot4>The man picks up a pen from the desk and begins writing in the notebook.</shot4>
</story>

Please write three new examples and output them in the same format as the example.
Don't be too similar to the written examples.