

Test-Time Noise Guided Adaptation for Realistic Autoregressive Video Generation

Dimitrios Karageorgiou^{1,2}, Symeon Papadopoulos¹, Ioannis Kompatsiaris¹,
and Efstratios Gavves²

¹ Information Technologies Institute, CERTH, Thessaloniki, Greece

² University of Amsterdam, Amsterdam, The Netherlands

{dkarageo,papadop,ikom}@iti.gr, {d.karageorgiou,e.gavves}@uva.nl

Abstract. Autoregressive video diffusion models have enabled the generation of arbitrarily long videos by removing conditioning on future frames, thus greatly improving computational efficiency. Yet, they suffer from error accumulation over time, as the denoised sequence gradually drifts away from the conditioning distribution seen during training. Recent advances attempt to reduce this error by anchoring each generated frame to the learned manifold of real ones. However, even when all generated individual frames lie close to the real manifold, there are trajectories which the model lacks sufficient knowledge to continue without exiting it, thus reaching a *terminal point*. To prevent the model from being trapped in terminal points, we start from the hypothesis that for well-modeled future trajectories the distribution of the predicted noise should match the one of the forward noising process. To enforce such a prior at test time, we introduce *Terminal points Avoidance through Noise Guided Optimization (TANGO)*, which uses the diffusion model as a critic of its own outputs, by predicting one step forward and requiring an isotropic Gaussian noise prediction. We use the deviation from this expected noise distribution to search for an alternative trajectory that does not lead to a terminal point. Our approach achieves a 3.1% absolute improvement on VBench over state-of-the-art, while reducing Fréchet Video Distance by 28.3% on average across 15s videos. Our code is available on <https://mever-team.github.io/tango>.

1 Introduction

Autoregressive video diffusion models [19, 31, 51, 57] have recently emerged as a paradigm to generate arbitrarily long videos. To achieve this, they employ causal attention structures that condition only on past frames to generate future ones – in practice on a limited number of them through rolling cache mechanisms [57]. Previously, bi-directional diffusion models [2, 47, 52] had to process all frames at once, while typically operating on fixed lengths considered during training. This was limiting the maximum attainable length, as the computational cost of commonly employed architectures, such as the UNets [39] and diffusion transformers [36], scales quadratically with respect to length.




Fig. 1: While terminal points reside in the learned real manifold, no continuation from them exists inside it. Thus, they act as traps that force the generated trajectory to continue in unrealistic regions. By employing TANGO, we avoid terminal points at test time through noise guided optimization, thus keeping the generative process within the real manifold. The upper half of the figure visualizes the trajectory correction performed by TANGO. As TANGO optimizes the visual trajectory of an autoregressive diffusion transformer by predicting a step forward, without relying on the initial condition, it can be uniformly applied for text, image and video-to-video generation. The bottom half of the figure displays three corresponding examples.

While such advances enabled the generation of long videos, autoregressive approaches suffer from a training-test-time mismatch in the distribution of the past visual sequences they attend to, as typically, they are trained on a set of past trajectories originating from real videos. Instead, at inference, they have to attend to their own generated frame sequences. As the generative process progresses, small shifts in the distribution of the generated sequences from the one seen during training accumulate, causing the autoregressive process to eventually collapse [19]. We distinguish two reasons that may cause the generated sequence of an autoregressive diffusion model to shift away from its training distribution. The first – which is the one discussed in recent literature – is related to the imperfect noise removal during the generation of individual future frames. The denoising procedure, due to its inherent aleatoric uncertainty, as well as due to imperfections during the modeling of the training distribution, fails to match the noise levels of the training data. Therefore, denoised frames end up outside of the learned manifold of real ones. Recently, several techniques for tackling this issue have emerged. They include the temporal unrolling of the generation process during training [19, 60], the repetition of the denoising timesteps during test time [21], or the incorporation of attention sinks [10, 31, 56] to anchor the generation manifold close to the initial frames. Yet, such approaches assume that the generated sequence will match the training distribution given all its individual frames reside in the learned manifold of real ones.

Instead, in this work we posit that even when all individual frames of a generated trajectory originate from the learned manifold of real samples, their joint distribution can differ from the one seen during training, i.e., the autoregressive process eventually generates a trajectory, for which, assuming a finite training set, no continuation exists under the learned manifold – thus, reaching a *terminal point*. Forcing the autoregressive process to continue causes it to exit the learned manifold of real samples and the produced video to become highly unrealistic. In this work, we formulate an approach to avoid terminal points and consequently, to prevent the model from generating sequences it does not know how to continue.

Understanding whether a point in the generated trajectory is terminal ultimately requires capturing the out-of-distribution conditioning sequences of the autoregressive diffusion model. To this end, our key proposition is that for well-modeled conditioning sequences, the distribution of the predicted noise should match the one considered during the forward noising process. We exploit this observation to use the denoising model itself as a critic of its past outputs, by observing whether the attributes of the predicted noise match the ones of an isotropic Gaussian distribution. To guide the diffusion model away from terminal points at test time, we introduce the Terminal points Avoidance through Noise Guided Optimization (TANGO), a method that searches for alternative trajectories that do not lead to terminal points. This is achieved by predicting one step forward and performing a constrained optimization in the neighborhood of the terminal point. Ultimately, the proposed solution continuously corrects the generated trajectory to avoid being trapped in points from which the generative model lacks the knowledge to continue. A summary of our work is presented in Fig. 1.

Our paper features the following key contributions:

- We highlight that terminal points prevent the use of many real trajectories as valid conditioning sequences for an autoregressive video generation model.
- We posit that it is possible to capture terminal points in autoregressive diffusion models, by observing the attributes of the predicted noise.
- We introduce TANGO, a test-time noise guided adaptation approach, to avoid generating trajectories that lead to terminal points, by enabling a diffusion model to criticize its own outputs.
- Through noise guided optimization we achieve a 3.1% absolute improvement over state-of-the-art on VBench [20], while reducing Fréchet Video Distance by 28.3% on average across 15 s videos.

2 Related Work

Video Diffusion Models. Video generative modeling has been revolutionized by the advent of diffusion models [9, 15], which have largely supplanted earlier adversarial [8, 44] and token-based autoregressive [46, 55] approaches. Initially built upon 3D U-Net architectures [16, 40] and subsequently transitioning towards Diffusion Transformers (DiTs) [36], these models excel at generating high-fidelity, temporally consistent frames. Foundational works primarily focused on generating fixed-length clips by extending image diffusion models with temporal attention or 3D convolutions [4, 16]. Scaling these architectures has led to powerful foundation models capable of simulating highly realistic physical world dynamics [1, 5]. Concurrently, models like Lumiere [3] have demonstrated the efficacy of Space-Time U-Nets (STUNets) to synthesize the full temporal duration in a single pass, contrasting with earlier cascaded temporal super-resolution methods. Recent state-of-the-art bi-directional diffusion models [2, 47] process the entire video sequence simultaneously, leveraging global temporal context to ensure high consistency across all frames. However, this global processing paradigm bounds the maximum sequence length that can be generated due to the quadratic computational complexity of attention mechanisms and the fixed context windows considered during training.

Autoregressive Video Generation. Autoregressive video generation has emerged as a prominent paradigm, framing video synthesis as a sequence continuation task [6, 11, 23]. This approach has been heavily driven by the development of highly compressive visual tokenizers [58] that allow foundation models to process multimodal inputs as discrete tokens [27]. To maintain temporal coherence across longer sequences, recent methods like StreamingT2V [14] utilize specialized short- and long-term memory blocks. Recently, Yin et al. [57] introduced an approach for distilling a bi-directional diffusion transformer into an autoregressive one with a causal attention structure and a rolling KV cache, ultimately formulating text-, image- and video-to-video synthesis under a single pipeline, depending on the number of frames used to initialize the cache. However, autoregressive generation introduces significant challenges, most notably exposure bias. Because the model is trained to predict the next frames conditioned on real video frames, relying




Fig. 2: Overview of the test-time adaptation process for avoiding terminal points through noise guidance. The current output of the denoiser is criticized by moving a step forward and inspecting the properties of the predicted noise residual. Deviations from an isotropic Gaussian distribution are captured by a noise consistency objective. This objective is employed under a constrained optimization process, along with a regularization objective, to update a trainable copy of the denoiser. The output of this updated denoiser becomes the final output of each step.

on its own imperfect generated frames during inference creates a mismatch with respect to its training distribution [19]. As generation progresses, minor deviations of the generated frame sequence from the training distribution compound, leading to severe error accumulation and the eventual collapse of the generated trajectory.

Limiting Out-of-Distribution Generated Outputs. One line of work attempts to mitigate this at training time by unrolling the generation process [19, 34, 60] to expose the model to its own predictions. However, this significantly increases training costs and memory overhead, while only mitigating error accumulation inside the fixed window considered during training. More efficiently, other works fine-tune noise enhancing networks to prevent the generation of low-fidelity frames [59]. Another prominent technique is the incorporation of attention sinks [10, 31, 56] to anchor the generative process to the initial frames, effectively narrowing the generation manifold. Alternatively, a growing body of tuning-free inference methods attempts to maintain consistency by rescheduling noise sequences for long-range correlation [37], applying temporal co-denoising across overlapping clips [48], or balancing multi-band frequency distributions during the denoising process to prevent high-frequency detail degradation [33]. Other approaches impose test-time objectives to align with human preference [24, 38], attempt to stabilize the output by repeating denoising timesteps during test time [21] so the generated frames converge to the learned real manifold or optimize the model itself to stick to the

anchors [50]. Common ground in such approaches is the attempt to bring each individual generated frame into the learned manifold of real frames. When applied in autoregressive setups, they assume that if a generative model is conditioned on individual frames that resemble its training samples, it will successfully generate trajectories in the same manifold. Yet, their joint distribution is not guaranteed to match the training one, eventually leading to highly unrealistic outputs.

In this paper, we argue that in order for an autoregressive model to keep generating realistic outputs over time, it is crucial for the joint distribution of its conditioning frames to not deviate from its training one, i.e. to avoid reaching terminal points. To this end, we introduce our noise guided formulation to enable an autoregressive model to capture terminal points, by inspecting the attributes of its predicted noise. Then, we employ noise guidance as a test-time optimization target to guide the diffusion model to generate trajectories without terminal points, significantly increasing the realism of the generated content, while allowing the model to explore trajectories throughout its entire learned manifold of real frames.

3 Terminal Points Avoidance Through Noise Guided Optimization

3.1 The Issue of Terminal Points in Generated Trajectories

Let $\mathcal{V} = \{x_1, x_2, \dots\}$ denote a video, where $x_i \in \mathbb{R}^{h \times w \times c}$ are its frames, and $p(\mathcal{V})$ the distribution of real videos. Under a typical autoregressive factorization [19, 57] the objective of a video generation model is to match the conditional distribution $p(x_i|x_{<i})$, i.e. to generate a plausible continuation frame given a past trajectory, as $x_i \in \text{supp}(p(x_i|x_{<i})) \Rightarrow \{x_{<i}, x_i\} \in \text{supp}(p(\mathcal{V}))$. As we consider the time horizon to be infinite, for every generated past trajectory there should be real future trajectories, thus $\emptyset \neq \text{supp}(p(x_{i+1}|\{x_{<i}, x_i\})) \subset \text{supp}(p(\mathcal{V}))$. In practice however, an autoregressive generative model is trained on a limited set of finite length trajectories $\mathcal{D} \sim p_{\mathcal{D}}(\mathcal{V})$, where $\text{supp}(p_{\mathcal{D}}(\mathcal{V})) \subset \text{supp}(p(\mathcal{V}))$, modeling the distribution $p_{\mathcal{D}}(x_i|x_{<i})$. In this latter case, the assumption of an infinite time horizon no longer holds, so $\exists v = \{x_{<i}, x_i\}, x_i \in \text{supp}(p_{\mathcal{D}}(x_i|x_{<i}))$ such that $\nexists u = \{v, x_{i+1}\} \in \text{supp}(p_{\mathcal{D}}(\mathcal{V})), x_{i+1} \in \text{supp}(p_{\mathcal{D}}(x_{i+1}|v))$. Consequently, even if an autoregressive generative model learns to generate perfectly valid individual future frames with respect to its training distribution, it will sooner or later generate frame sequences that jointly constitute out-of-distribution samples with respect to it. Therefore, it will be trapped in regions of the learned manifold where no valid continuation exists without exiting it. We call these regions *terminal points*. A visualization of a terminal point is provided in the upper half of Fig. 1.

In order for an autoregressive model to avoid generating trajectories with terminal points it requires both a mechanism to identify them and an approach to guide it away from them. In this work we contribute to both directions. An overview of the proposed solution is provided in Fig. 2.

3.2 Noise Consistency Hypothesis for Capturing Terminal Points

From now on, we assume that the autoregressive video generation model is a diffusion one, and, for simplicity of notation and without loss of generality, we base our derivation on the score matching formulation [42]. Derivations for other formulations, like DDIM [41] or flow matching [30], can be directly constructed from it – we provide a derivation for the latter in the supplementary material.

Diffusion models are fundamentally based on the idea of moving from any random point in space along the direction that maximizes the likelihood of a sample being real, eventually reaching a point in the manifold of real samples. It has been shown [25, 42] that a model can learn to approximate this direction, starting from any point in space, by augmenting a set of real samples x^0 with isotropic Gaussian noise during the forward diffusion process:

$$x^t = \alpha^t x^0 + \sigma^t \epsilon, \epsilon \sim \mathcal{N}(0, I) \quad (1)$$

where α^t and σ^t are scalars that control the noise level according to a noise schedule. t defines the denoising timesteps indicated by the latter, and should not be confused with the temporal dimensionality of the generated video – denoted by i in the current context. Then, learning a model ϵ_θ to remove this noise, where θ is its set of trainable parameters, ultimately approximates the score function:

$$s(x^t, t) = \nabla_{x^t} \log p(x^t) \approx -\frac{\epsilon_\theta(x^t, t)}{\sigma_t} = s_\theta(x^t, t) \quad (2)$$

The denoiser is typically trained to minimize the mean squared error with respect to the added noise:

$$\mathcal{L}(\theta) = \mathbb{E}_{t, x^0, \epsilon} \|\epsilon_\theta(x^t, t) - \epsilon\|_2^2 \quad (3)$$

Under a global video denoising procedure, i.e. when $x^0 = \mathcal{V} \sim p_{\mathcal{D}}(\mathcal{V})$, the only cause of $\mathcal{L}(\theta) > 0$ during testing can be the imperfect modeling of the score function due to the unavoidably limited noise samples considered during training and the limited capacity of the model. However, under an autoregressive factorization, the score function takes the following form:

$$s(x_i^t, t, x_{<i}^0) = \nabla_{x^t} \log p(x_i^t | x_{<i}^0) \approx -\frac{\epsilon_\theta(x_i^t, t, x_{<i}^0)}{\sigma_t} = s_\theta(x_i^t, t, x_{<i}^0) \quad (4)$$

Therefore, any sequence $x_{<i}^0 \notin \text{supp}(p_{\mathcal{D}}(\mathcal{V}))$ is not considered during the minimization of $\mathcal{L}(\theta)$, even if we could assume an infinite amount of noise samples and a model of adequate capacity. Thus, the output of ϵ_θ is undefined in these regions and is allowed to violate the properties of isotropic Gaussian noise assumed in Eq. (1). Overall, we exploit such violations in order to use ϵ_θ as a critic of its own past outputs.

We posit that at a denoising timestep t it is possible to examine whether a predicted $\hat{x}_i^0$ constitutes a terminal point by examining whether the predicted look-ahead residuals $\hat{r}_{i+1} = \epsilon_\theta(x_{i+1}^t, t, \{x_{<i}, \hat{x}_i^0\})$ violate the properties of isotropic Gaussian noise.

3.3 Avoiding Terminal Points by Test-Time Adaptation

To capture deviations of $\hat{r}_{i+1} \in \mathbb{R}^{h \times w \times c}$ from isotropic Gaussian noise we compute, over its pixels, the distance of its statistical moments from the defining ones of $\mathcal{N}(0, I)$ in the spatial domain, as well as its flatness in the frequency domain.

First and Second Moments: Let $\mu_{\hat{r}_{i+1}}$ and $\sigma_{\hat{r}_{i+1}}^2$ be the empirical mean and biased variance of $\hat{r}_{i+1}$. We penalize deviations from the defining zero mean and unity variance of an isotropic Gaussian distribution, formulating the loss:

$$\mathcal{L}_a = \lambda_\mu (\mu_{\hat{r}_{i+1}})^2 + \lambda_\sigma (\sigma_{\hat{r}_{i+1}}^2 - 1)^2 \quad (5)$$

where λ_μ and λ_σ are scalar hyperparameters.

Higher-Order Moments: To detect asymmetry and structural artifacts (heavy tails), we use the z -scores $z = (\hat{r}_{i+1} - \mu_{\hat{r}_{i+1}}) / \sqrt{\sigma_{\hat{r}_{i+1}}^2}$ to compute skewness and excess kurtosis. We again penalize their deviations from the values they take under an isotropic Gaussian distribution, i.e., zero:

$$\mathcal{L}_b = \lambda_\gamma \underbrace{\mu(z^3)^2}_{\text{skewness}} + \lambda_\kappa \underbrace{(\mu(z^4) - 3)^2}_{\text{excess kurtosis}} \quad (6)$$

where $\mu(\cdot)$ is a pixel-wise mean operation and $\lambda_\gamma, \lambda_\kappa$ are scalar hyperparameters.

Spectral Flatness: One key property of isotropic noise is its equal energy levels in all frequencies, thus a flat power spectral density [35]. A non-flat power spectral density would indicate spatial correlation, thus the denoiser would also be removing visual content along with the added noise. To find out whether the predicted noise residual is isotropic we evaluate the spectral flatness measure (SFM) [13]. Let $S(f) = |\mathcal{F}(\hat{r}_{i+1})|^2$ be its power spectral density, where $\mathcal{F}$ denotes the Discrete Fourier Transform. To ensure numerical stability and avoid log-zero errors given a zero-mean signal, we exclude the DC component ($f = 0$) and add a small constant ζ . Then, we define the spectral flatness loss that penalizes deviations from 1.0 (perfect isotropic noise):

$$\mathcal{L}_s = \lambda_s \left(1 - \frac{\exp\left(\frac{1}{|F|} \sum_{f \in F} \ln(S(f) + \zeta)\right)}{\frac{1}{|F|} \sum_{f \in F} (S(f) + \zeta)} \right)^2, \quad (7)$$

where F is the set of non-zero frequencies and λ_s a scalar hyperparameter.

Moments in Low Frequencies: As $\hat{r}_{i+1}$ is dominated by high frequencies, the computation of statistical moments over it is mostly overwhelmed by them. However, the energy of real visual content mostly resides in low frequencies [12]. To explicitly enforce Gaussian properties in low frequencies, we also compute skewness and excess kurtosis over a low-pass filtered variant of $\hat{r}_{i+1}$, whose z -scores are denoted z^l . We define the corresponding loss as follows:

$$\mathcal{L}_l = \lambda_{l\mu} (\mu((z^l)^3))^2 + \lambda_{l\sigma} (\mu((z^l)^4) - 3)^2 \quad (8)$$

where $\lambda_{l\mu}$, $\lambda_{l\sigma}$ are once again scalar hyperparameters. Then, the total noise consistency objective becomes:

$$\mathcal{L}_{\mathcal{N}} = \mathcal{L}_a + \mathcal{L}_b + \mathcal{L}_s + \mathcal{L}_l \quad (9)$$

To enforce this consistency at test time, we introduce a set of adaptive parameters ϕ , initialized with the pre-trained weights θ . At each autoregressive step i , we optimize ϕ to generate $\hat{x}_i^0(\phi)$ such that it satisfies $\mathcal{L}_{\mathcal{N}}$ computed using the frozen critic ϵ_{θ} .

To avoid trivial solutions, we formulate the adaptation as a constrained optimization problem. We update ϕ to minimize $\mathcal{L}_{\mathcal{N}}$ subject to a regularization constraint that anchors the search process in the neighborhood of $\hat{x}_i^0(\theta)$, which denotes the original model’s prediction. The total objective becomes:

$$\mathcal{J}(\phi) = \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)}[\mathcal{L}_{\mathcal{N}}] + \lambda_{\text{reg}} \|\hat{x}_i^0(\phi) - \hat{x}_i^0(\theta)\|_2^2 \quad (10)$$

where λ_{reg} denotes a weighting hyperparameter. In practice, we approximate $\mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)}[\mathcal{L}_{\mathcal{N}}]$ via a batch of B noise vectors at each optimization step, controllable for the available compute power at test time. By minimizing Eq. (10), the optimized model is forced to avoid terminal points in the frozen manifold.

4 Experiments

4.1 Implementation Details

To facilitate our experiments on noise guided optimization, we follow the recent literature [31, 56] and initialize our denoiser ϵ_{θ} according to the distribution matching distillation scheme with self-unrolling, causal attention and KV-cache introduced by Huang et al. [19], starting from the bi-directional video diffusion transformer [16] of Wan2.1-T2V-1.3B [47] and using the inference optimizations of [34]. As the employed model is based on flow matching [30], we solve the equation of the predicted flow with respect to noise in order to obtain the predicted noise residual $\hat{r}_{i+1}$ that we employ to guide the denoising procedure. The derivation is provided in the supplementary material. The RGB resolution of the generated video is 832×480 at 16FPS and with an unlimited time horizon, due to a rolling KV cache that retains 21 latent frames (≈ 5 s of video). All the computations are taking place in the compressed latent space of a 3D variational autoencoder [26], that provides $\times 8$ spatial compression and $\times 4$ temporal compression, with the exception of the initial frame that is not temporally compressed to capture richer context. Thus, $h = 60$, $w = 104$, while $c = 16$ is used for the latent channels. During each optimization stage, we introduce a small set of trainable parameters to the self-attention layers of the causal diffusion transformer through a low-rank adaptation [17] of rank 8. We set $\lambda_{\mu} = \lambda_{\sigma} = \lambda_{\gamma} = \lambda_{\kappa} = \lambda_s = 0.1$, $\lambda_{l\mu} = \lambda_{l\sigma} = 0.2$ and $\lambda_{\text{reg}} = 0.9$ based on a preliminary tuning on the test set of MSR-VTT [53], which we do not employ further in our evaluation, to avoid any potential overfitting. We optimize the

model with a learning rate of 2×10^{-6} for a total of $B = 10$ noise vectors, across a single epoch with batch size 2, using the AdamW [32] optimizer. Experiments were conducted on a single Nvidia RTX Pro 6000 GPU with 96 GB of memory.

4.2 Benchmarks & Considered Generative Approaches

To compute the realism of generated videos we follow recent literature [19, 31, 57] and employ the VBench [20] suite. It defines a set of 946 textual prompts, spanning 11 evaluation dimensions and 8 categories of visual content, which are employed to compute a set of 16 faithfully crafted metrics. The latter are used to compute a quality and semantic score, as well as a single total score, indicative of the overall performance of a generative approach. We use all of them under a text-to-video generation pipeline. Moreover, it is not possible to account, solely through handcrafted metrics, for every possible aspect that contributes to the visual faithfulness of a video. For this reason, we also employ Fréchet Video Distance (FVD) [45], to compute the distance of the distribution of videos produced by a generative model, to the one of real videos. As a reference dataset of real videos we employ the LV-Bench [60], which combines videos from DanceTrack [43], GOT-10k [18], HD-VILA-100M [54] and ShareGPT4V [7], considering human, animal and environment-focused visual scenes, along with curated captions. To better focus on the visuals of the video, limiting the effects of a potential contextual divergence between the generated and real videos, we employ them under a video-to-video continuation pipeline.

In our experiments, we consider a set of recently proposed autoregressive video generation approaches, with publicly available code and of similar scale. In particular, we consider CausVid [57], Self-Forcing [19], RollingForcing [31], LongLive [56] and RewardForcing [34], as all of them are initialized with the Wan2.1-T2V-1.3B [47] bi-directional video diffusion transformer, thus eliminating, as much as possible, the effects of different training data and scales, enabling us to better focus on the algorithms. For a fair comparison, we use the public code of each approach to generate videos of equal length (15 seconds).

4.3 Comparison against State-of-the-Art

We compute all VBench metrics for all considered approaches and present them in Tab. 1. Overall, we observe an absolute performance increase of 4.8% in semantic score, 2.6% in quality score and 3.1% in the total score, as well as the best performance in most individual metrics with respect to the previous state-of-the-art. Noticeably, even in the dynamic degree metric that quantifies the diversity of generated content, our approach is only second to CausVid [57], which does so at the cost of semantic coherence in the generated frames. Instead, approaches that primarily rely on anchoring to the early frames to maintain visual consistency, like [31, 56], compromise their ability to generate diverse content.

Table 1: Comparison against state-of-the-art autoregressive video generation approaches on 16 visual aspects. By avoiding terminal points at test-time, TANGO enhances visual quality and semantic consistency, while better preserving content diversity. Best values are highlighted in bold, while second best are underlined.

Approach	Aesthetic Quality	Appearance Style	Background Consistency	Color	Dynamic Degree	Human Action	Imaging Quality	Motion Smoothness	Multiple Objects	Object Class
RollingForcing [31]	0.658	0.205	0.957	0.819	0.403	0.940	<u>0.712</u>	<u>0.986</u>	<u>0.932</u>	0.935
SelfForcing [19]	0.644	0.206	0.942	0.848	0.611	0.950	0.696	0.985	0.896	0.959
CausVid [57]	<u>0.665</u>	<u>0.242</u>	0.966	0.824	0.872	0.998	0.702	0.976	0.747	0.937
RewardForcing [34]	0.650	0.205	0.954	0.896	0.625	<u>0.980</u>	0.687	0.984	0.845	0.940
LongLive [56]	0.660	0.205	0.958	<u>0.885</u>	0.361	<u>0.960</u>	0.695	0.988	0.851	<u>0.969</u>
TANGO (ours)	0.690	0.257	<u>0.962</u>	0.850	<u>0.769</u>	0.960	0.727	0.985	0.937	0.972

Approach	Overall Consistency	Scene	Spatial Relationship	Subject Consistency	Temporal Flickering	Temporal Style	Quality Score	Semantic Score	Total
RollingForcing [31]	0.269	0.563	0.747	0.958	0.991	0.241	0.802	0.814	0.804
SelfForcing [19]	0.271	0.578	0.768	0.928	0.989	0.249	0.832	0.808	0.827
CausVid [57]	<u>0.275</u>	0.556	0.667	0.933	0.938	<u>0.252</u>	0.837	0.794	0.828
RewardForcing [34]	0.269	0.539	0.830	0.932	0.986	0.244	0.835	0.804	0.829
LongLive [56]	0.268	<u>0.591</u>	0.801	<u>0.953</u>	0.992	0.242	<u>0.838</u>	0.812	<u>0.832</u>
TANGO (ours)	0.283	0.603	0.859	0.945	0.989	0.275	0.864	0.860	0.863

Table 2: Ablation studies of the key losses of noise guided optimization. Average Fréchet Video Distance, between ground-truth and generated video segments, is reported over 3-second segments as well as over the entire videos. Lower is better. Best values in bold.

Ablation (FVD ↓)	1-3s	4-6s	7-9s	10-12s	13-15s	Total
TANGO (ours)	365.1	408.8	412.7	441.4	438.8	413.3
w/o $\mathcal{L}_a$	364.6	418.1	467.1	469.8	488.2	441.6
w/o $\mathcal{L}_l$	380.3	452.8	518.4	534.8	545.2	486.3
w/o $\mathcal{L}_s$	369.3	443.9	507.1	530.6	578.4	487.9
w/o $\mathcal{L}_b$	363.7	449.8	561.2	587.1	572.3	506.8
w/o $\mathcal{L}_{reg}$	420.1	634.4	630.9	823.9	811.7	664.2
w/o $\mathcal{L}_N$ (baseline)	367.9	554.6	600.7	645.4	715.6	576.8

4.4 Ablation Studies

Loss terms: Since our objective is to prevent the autoregressive process from exiting the learned manifold of real videos, and not to merely enforce individual visual attributes, we ablate our key choices by directly measuring the divergence of the distribution of generated videos from the one of real clips. To this end, we compute the average FVD for consecutive 3-second segments of the generated videos from their corresponding segments in the real videos of LV-Bench [60]. We report results in Tab. 2. We observe that higher-order moments are very important for capturing deviations in the predicted noise residuals from the distribution of isotropic Gaussian noise. Removing the corresponding loss impacts performance the most compared to the other losses related to noise consistency. Next, we observe that removing the objective of spectral flatness causes a significant performance drop, validating our hypothesis that spatial correlations would remove actual visual content, increasing the drift with respect to the considered distribution of real videos. Similarly, explicitly computing moments in low fre-

Table 3: Hyperparameters ablation studies. Average Fréchet Video Distance, between ground-truth and generated segments, is reported for 15-second videos. Lower is better.

LoRA R.	FVD ↓	B	FVD ↓	λ_{reg}	FVD ↓
2	596.2	1	473.7	0.1	656.5
4	496.3	5	461.8	0.5	449.3
8	413.3	10	413.3	0.9	413.3
16	425.6	20	399.4	10	475.6

(a) Ablations of key hyperparameters involved in test-time adaptation.

λ_μ	FVD ↓	λ_σ	FVD ↓	λ_{l_μ}	FVD ↓	λ_{l_σ}	FVD ↓	λ_γ	FVD ↓	λ_κ	FVD ↓	λ_s	FVD ↓
.01	503.8	.01	482.8	.02	419.8	.02	482.8	.01	468.1	.01	479	.01	477.4
.1	413.3	.1	413.3	.2	413.3	.2	413.3	.1	413.3	.1	413.3	.1	413.3
1	449.3	1	540.7	2	464.7	2	464.7	1	442.7	1	482.6	1	429.1

(b) Ablations of the weights of the noise consistency objective.**Table 4:** Runtime and memory usage comparison of the proposed test-time approach with ones that enhance performance by employing a larger model.

Model	Params	VBench (↑)	Runtime (↓)	Memory (↓)
Baseline (SelfForcing) [19]	1.3B	0.827	1.2 min	23.4 GiB
HunyuanVideo [28]	13B	0.832	92.3 min	62.5 GiB
SVI-1.0 (Wan2.1) [29]	14B	0.837	30.6 min	37.9 GiB
SVI-2.0 (Wan2.1) [29]	14B	0.849	39.8 min	37.9 GiB
SVI-2.0-Pro (Wan2.2) [29]	27B	0.869	43.2 min	68.4 GiB
TANGO (ours)	1.3B	0.863	8.6 min	37.6 GiB

quencies helps the optimization process better capture visible inconsistencies. Also, results show that regularization is crucial for preventing the optimization process from converging to trivial solutions. In fact, performing optimization without regularization degrades performance over the baseline. Overall, compared to the baseline, our test-time optimization process reduces FVD by 38.7% after 15s of generation, while also reducing it by 28.3% on average at a video level. This indicates that our approach not only helps improve specific aspects of the generated videos, as previously presented in Tab. 1, but more generally, reduces the distance between the distributions of the generated and the real videos.

Hyperparameters: We further ablate the key hyperparameters and report the results in Tab. 3. Overall, while very low LoRA ranks could even worsen performance, gains from an increase in trainable parameters eventually diminish. Instead, we observe that further increasing the number of noise vectors B provides additional performance. Yet, as it comes at the cost of computational power, we set $B = 10$ as a reasonable trade-off between efficiency and performance. As for the regularization coefficient and the weights of our noise consistency objective, we observe higher performance when they range in the same order of magnitude.

4.5 Runtime & Memory Usage Analysis

We compared the runtime and memory usage of our test-time optimization approach against open-source video generation methods [28, 29] that increase visual




Fig. 3: Qualitative comparison against state-of-the-art approaches.




Fig. 4: Failure cases. From left to right, frames indicative of four temporal video segments, of 15s total duration, are presented. Two videos are considered, spanning the niche topics of exotic physics and nanobiology. While results are visually pleasant, technically, they are completely inaccurate.

performance primarily by scaling up the model’s size. We report the respective results in Tab. 4. Overall, steering at test-time a smaller autoregressive video diffusion model [19] away from terminal points helps increase visual performance, while maintaining better trade-offs in runtime and memory usage.

4.6 Qualitative Evaluation

In Fig. 3 we present a qualitative evaluation against the two best performing approaches as indicated by Tab. 1. We consider indicative frames from four different segments of 15-second videos. As shown in the first two videos, in complex scenes or scenes generated based on ambiguous textual prompts, generative models quickly get out of the learned real manifold, producing unrealistic scenes with visible artifacts. Restricting the exploration space in a narrow region around the initial points of the generated trajectory on several occasions improves realism at the cost of dynamic degree. Instead, employing TANGO to avoid being trapped in terminal points improves realism by finding a valid continuation under the limited world knowledge of the diffusion model.

4.7 Discussion & Failure Cases

Any test-time adaptation approach effectively allocates inference time compute power to find a better solution under a modeled distribution, enabling the model to move away from local maxima where it may get stuck otherwise. Our approach is not an exception in this. However, as the field of large language models has shown, enabling a model to reason about its outputs [49] and fix

its mistakes is a crucial component of the inference process, with the potential of enabling models to extrapolate beyond their training regime. However, in contrast to language-based autoregressive modeling, autoregressive video diffusion models lacked a formulation to move them away from terminal points and in this work we introduce such an approach. Future work could span architectures that inherently learn to quantify their epistemic uncertainty [22] without requiring gradient-descent-based optimization procedures.

Given a fixed scale, our test-time optimization approach enhances the capabilities of a model. Yet, finding a better solution under its learned manifold requires the existence of one. For visual topics with significant underrepresentation in the training data, finding a trajectory without terminal points is challenging. Case in point, Fig. 4 displays the frames of two videos, with the topic of their prompts spanning exotic physics and nanobiology. While noise guidance finds trajectories that look pleasant to the eye, from a technical perspective the results are utterly wrong. Therefore, test-time adaptation is not competitive, but complementary to scaling up training data and model capacity, as the latter is the one that increases the space of valid solutions. Instead, noise guidance attempts to find a better one.

5 Conclusion

In this work we highlighted that it is not enough for an autoregressive video generation model to merely generate future trajectories that reside in the learned real manifold, as several points in it constitute terminal points. Once reached, no continuation exists without exiting it. To avoid reaching such points, and therefore to avoid the generation of unrealistic visual content, we introduced a noise guided test-time adaptation mechanism, that uses the diffusion model itself as a critic of its own outputs. It does so by assuming that when attempting to continue a trajectory that reached a terminal point, the noise properties assumed during the forward noising process will no longer hold for the predictions of the denoiser. We exploited this observation in a constrained optimization process to move away from terminal points. Overall, by employing this test-time approach, we managed to improve state-of-the-art performance across several metrics of the realism of videos, without requiring any additional data.

Limitations: Any test-time adaptation approach effectively searches for an improved solution under the manifold learned during pre-training and will eventually fail if no such solution exists. Therefore, we view data and model scaling as complementary directions for increasing the space of possible solutions. Yet, our proposed noise guided optimization process contributes to finding such better solutions under the learned manifold of an autoregressive video diffusion model.

Acknowledgements

This work was supported by the Horizon Europe projects ELIAS (grant no. 101120237) and ELLIOT (grant no. 101214398). The computational resources were granted with the support of GRNET.

References

1. Agarwal, N., Ali, A., Bala, M., Balaji, Y., Barker, E., Cai, T., Chattopadhyay, P., Chen, Y., Cui, Y., Ding, Y., et al.: Cosmos world foundation model platform for physical ai. arXiv preprint arXiv:2501.03575 (2025)
2. Ali, A., Bai, J., Bala, M., Balaji, Y., Blakeman, A., Cai, T., Cao, J., Cao, T., Cha, E., Chao, Y.W., et al.: World simulation with video foundation models for physical ai. arXiv preprint arXiv:2511.00062 (2025)
3. Bar-Tal, O., Chefer, H., Tov, O., Herrmann, C., Paiss, R., Zada, S., Ephrat, A., Hur, J., Liu, G., Raj, A., et al.: Lumiere: A space-time diffusion model for video generation. In: SIGGRAPH Asia 2024 Conference Papers. pp. 1–11 (2024)
4. Blattmann, A., Rombach, R., Ling, H., Dockhorn, T., Kim, S.W., Fidler, S., Kreis, K.: Align your latents: High-resolution video synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 22563–22575 (2023)
5. Brooks, T., Peebles, B., Holmes, C., DePue, W., Guo, Y., Jing, L., Schnurr, D., Taylor, J., Luhman, T., Luhman, E., Ng, C., Wang, R., Ramesh, A.: Video generation models as world simulators (2024), <https://openai.com/index/video-generation-models-as-world-simulators/>
6. Chen, G., Lin, D., Yang, J., Lin, C., Zhu, J., Fan, M., Zhang, H., Chen, S., Chen, Z., Ma, C., et al.: Skyreels-v2: Infinite-length film generative model. arXiv preprint arXiv:2504.13074 (2025)
7. Chen, L., Li, J., Dong, X., Zhang, P., He, C., Wang, J., Zhao, F., Lin, D.: Sharegpt4v: Improving large multi-modal models with better captions. In: European Conference on Computer Vision. pp. 370–387. Springer (2024)
8. Clark, A., Donahue, J., Simonyan, K.: Adversarial video generation on complex datasets. arXiv preprint arXiv:1907.06571 (2019)
9. Croitoru, F.A., Hondru, V., Ionescu, R.T., Shah, M.: Diffusion models in vision: A survey. IEEE transactions on pattern analysis and machine intelligence **45**(9), 10850–10869 (2023)
10. Cui, J., Wu, J., Li, M., Yang, T., Li, X., Wang, R., Bai, A., Ban, Y., Hsieh, C.J.: Lol: Longer than longer, scaling video generation to hour. arXiv preprint arXiv:2601.16914 (2026)
11. Deng, H., Pan, T., Diao, H., Luo, Z., Cui, Y., Lu, H., Shan, S., Qi, Y., Wang, X.: Autoregressive video generation without vector quantization. In: International Conference on Learning Representations (ICLR) (2025)
12. Field, D.J.: Relations between the statistics of natural images and the response properties of cortical cells. Journal of the Optical Society of America A **4**(12), 2379–2394 (1987)
13. Gray, A., Markel, J.: A spectral-flatness measure for studying the autocorrelation method of linear prediction of speech analysis. IEEE Transactions on Acoustics, Speech, and Signal Processing **22**(3), 207–217 (1974)
14. Henschel, R., Khachatryan, L., Poghosyan, H., Hayrapetyan, D., Tadevosyan, V., Wang, Z., Navasardyan, S., Shi, H.: Streamingt2v: Consistent, dynamic, and extendable long video generation from text. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 2568–2577 (2025)
15. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020)
16. Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. Advances in neural information processing systems **35**, 8633–8646 (2022)

17. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. *International Conference on Learning Representations (ICLR)* (2022)
18. Huang, L., Zhao, X., Huang, K.: Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *IEEE transactions on pattern analysis and machine intelligence* **43**(5), 1562–1577 (2019)
19. Huang, X., Li, Z., He, G., Zhou, M., Shechtman, E.: Self forcing: Bridging the train-test gap in autoregressive video diffusion. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2025)
20. Huang, Z., He, Y., Yu, J., Zhang, F., Si, C., Jiang, Y., Zhang, Y., Wu, T., Jin, Q., Chanpaisit, N., et al.: Vbench: Comprehensive benchmark suite for video generative models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 21807–21818 (2024)
21. Jang, S., Ki, T., Jo, J., Xie, S., Yoon, J., Hwang, S.J.: Self-refining video sampling. In: *International Conference on Machine Learning (ICML)* (2026)
22. Jazbec, M., Wong-Toi, E., Xia, G., Zhang, D., Nalisnick, E., Mandt, S.: Generative uncertainty in diffusion models. In: *Uncertainty in Artificial Intelligence (UAI)* (2025)
23. Kim, J., Kang, J., Choi, J., Han, B.: Fifo-diffusion: Generating infinite videos from text without training. *Advances in Neural Information Processing Systems* **37**, 89834–89868 (2024)
24. Kim, S., Kim, M., Park, D.: Test-time alignment of diffusion models without reward over-optimization. In: *International Conference on Learning Representations (ICLR)* (2025)
25. Kingma, D., Salimans, T., Poole, B., Ho, J.: Variational diffusion models. *Advances in neural information processing systems* **34**, 21696–21707 (2021)
26. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: *International Conference on Learning Representations (ICLR)* (2014)
27. Kondratyuk, D., Yu, L., Gu, X., Lezama, J., Huang, J., Schindler, G., Hornung, R., Birodkar, V., Yan, J., Chiu, M.C., et al.: Videopoet: A large language model for zero-shot video generation. In: *International Conference on Machine Learning (ICML)* (2024)
28. Kong, W., Tian, Q., Zhang, Z., Min, R., Dai, Z., Zhou, J., Xiong, J., Li, X., Wu, B., Zhang, J., et al.: Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603* (2024)
29. Li, W., Pan, W., Luan, P.C., Gao, Y., Alahi, A.: Stable video infinity: Infinite-length video generation with error recycling. *arXiv preprint arXiv:2510.09212* (2025)
30. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: *International Conference on Learning Representations (ICLR)* (2023)
31. Liu, K., Hu, W., Xu, J., Shan, Y., Lu, S.: Rolling forcing: Autoregressive long video diffusion in real time. In: *International Conference on Learning Representations (ICLR)* (2026)
32. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: *International Conference on Learning Representations (ICLR)* (2019)
33. Lu, Y., Liang, Y., Zhu, L., Yang, Y.: Freelong: Training-free long video generation with spectralblend temporal attention. *Advances in Neural Information Processing Systems* **37**, 131434–131455 (2024)
34. Lu, Y., Zeng, Y., Li, H., Ouyang, H., Wang, Q., Cheng, K.L., Zhu, J., Cao, H., Zhang, Z., Zhu, X., et al.: Reward forcing: Efficient streaming video generation

- with rewarded distribution matching distillation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 34385–34397 (2026)
35. Madhu, N.: Note on measures for spectral flatness. *Electronics letters* **45**(23), 1195–1196 (2009)
 36. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4195–4205 (2023)
 37. Qiu, H., Xia, M., Zhang, Y., He, Y., Wang, X., Shan, Y., Liu, Z.: Freenoise: Tuning-free longer video diffusion via noise rescheduling. *arXiv preprint arXiv:2310.15169* (2023)
 38. Qu, L., Wang, Z., Zheng, N., Wang, W., Nie, L., Chua, T.S.: Ttom: Test-time optimization and memorization for compositional video generation. *arXiv preprint arXiv:2510.07940* (2025)
 39. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 234–241. Springer (2015)
 40. Singer, U., Polyak, A., Hayes, T., Yin, X., An, J., Zhang, S., Hu, Q., Yang, H., Ashual, O., Gafni, O., et al.: Make-a-video: Text-to-video generation without text-video data. In: *International Conference on Learning Representations (ICLR)* (2023)
 41. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: *International Conference on Learning Representations (ICLR)* (2021)
 42. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: *International Conference on Learning Representations (ICLR)* (2021)
 43. Sun, P., Cao, J., Jiang, Y., Yuan, Z., Bai, S., Kitani, K., Luo, P.: Dancetrack: Multi-object tracking in uniform appearance and diverse motion. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 20993–21002 (2022)
 44. Tulyakov, S., Liu, M.Y., Yang, X., Kautz, J.: Mocogan: Decomposing motion and content for video generation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1526–1535 (2018)
 45. Unterthiner, T., Van Steenkiste, S., Kurach, K., Marinier, R., Michalski, M., Gelly, S.: Fvd: A new metric for video generation. In: *Deep Generative Models for Highly Structured Data, ICLR 2019 Workshop* (2019)
 46. Villegas, R., Babaeizadeh, M., Kindermans, P.J., Moraldo, H., Zhang, H., Saffar, M.T., Castro, S., Kunze, J., Erhan, D.: Phenaki: Variable length video generation from open domain textual description. In: *International Conference on Learning Representations (ICLR)* (2023)
 47. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., et al.: Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314* (2025)
 48. Wang, F.Y., Chen, W., Song, G., Ye, H.J., Liu, Y., Li, H.: Gen-l-video: Multi-text to long video generation via temporal co-denoising. *arXiv preprint arXiv:2305.18264* (2023)
 49. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* **35**, 24824–24837 (2022)
 50. Xiang, X., Duan, Z., Zhang, G., Zhang, H., Gao, Z., Wu, J., Zhang, S., Wang, T., Fan, Q., Guo, C.: Pathwise test-time correction for autoregressive long video generation. *arXiv preprint arXiv:2602.05871* (2026)

51. Xie, D., Xu, Z., Hong, Y., Tan, H., Liu, D., Liu, F., Kaufman, A., Zhou, Y.: Progressive autoregressive video diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2025)
52. Xing, Z., Feng, Q., Chen, H., Dai, Q., Hu, H., Xu, H., Wu, Z., Jiang, Y.G.: A survey on video diffusion models. *ACM Comput. Surv.* **57**(2) (Nov 2024). <https://doi.org/10.1145/3696415>, <https://doi.org/10.1145/3696415>
53. Xu, J., Mei, T., Yao, T., Rui, Y.: Msr-vtt: A large video description dataset for bridging video and language. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5288–5296 (2016)
54. Xue, H., Hang, T., Zeng, Y., Sun, Y., Liu, B., Yang, H., Fu, J., Guo, B.: Advancing high-resolution video-language representation with large-scale video transcriptions. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5036–5045 (2022)
55. Yan, W., Zhang, Y., Abbeel, P., Srinivas, A.: Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157* (2021)
56. Yang, S., Huang, W., Chu, R., Xiao, Y., Zhao, Y., Wang, X., Li, M., Xie, E., Chen, Y., Lu, Y., et al.: Longlive: Real-time interactive long video generation. In: International Conference on Learning Representations (ICLR) (2026)
57. Yin, T., Zhang, Q., Zhang, R., Freeman, W.T., Durand, F., Shechtman, E., Huang, X.: From slow bidirectional to fast autoregressive video diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22963–22974 (2025)
58. Yu, L., Cheng, Y., Sohn, K., Lezama, J., Zhang, H., Chang, H., Hauptmann, A.G., Yang, M.H., Hao, Y., Essa, I., et al.: Magvit: Masked generative video transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10459–10469 (2023)
59. Yu, Z., Hayakawa, A., Ishii, M., Yu, Q., Shibuya, T., Zhang, J., Mitsufuji, Y.: Autorefiner: Improving autoregressive video diffusion models via reflective refinement over the stochastic sampling path. *arXiv preprint arXiv:2512.11203* (2025)
60. Zhang, Z., Chang, S., He, Y., Han, Y., Tang, J., Wang, F., Zhuang, B.: Blockvid: Block diffusion for high-quality and consistent minute-long video generation. *arXiv preprint arXiv:2511.22973* (2025)