

VIGOR: Video Geometry-Oriented Reward for Temporal Generative Alignment

Tengjiao Yin¹, Jinglei Shi^{1†}, Heng Guo², and Xi Wang³

¹ VCIP & TMCC & DISSec, College of Computer Science, Nankai University

² Beijing University of Posts and Telecommunications

³ LIX, École Polytechnique, IP Paris

<https://vigor-geometry-reward.com/>

Abstract. Video diffusion models lack explicit geometric supervision during training, leading to inconsistency artifacts such as object deformation, spatial drift, and depth violations in generated videos. To address this limitation, we propose a geometry-based reward model that leverages pretrained geometric foundation models to evaluate multi-view consistency through cross-frame reprojection error. Unlike previous geometric metrics that measure inconsistency in pixel space, where pixel intensity may introduce additional noise, our approach conducts error computation in a pointwise fashion, yielding a more physically grounded and robust error metric. Furthermore, we introduce a geometry-aware sampling strategy that filters out low-texture and non-semantic regions, focusing evaluation on geometrically meaningful areas with reliable correspondences to improve robustness. We apply this reward model to align video diffusion models through two complementary pathways: post-training of a bidirectional model via SFT or Reinforcement Learning and inference-time optimization of a Causal Video Model (e.g., Streaming video generator) via test-time scaling with our reward as a path verifier. Experimental results validate the effectiveness of our design, demonstrating that our geometry-based reward provides superior robustness compared to other variants. By enabling efficient inference-time scaling, our method offers a practical solution for enhancing open-source video models without requiring extensive computational resources for retraining.

Keywords: Video Diffusion Model · Geometry-based Reward · Multi-view Consistency · Causal Inference

1 Introduction

Video diffusion models have achieved remarkable progress in recent years [4, 14, 25, 37, 41, 49] demonstrating increasingly photorealistic and temporally coherent video synthesis. Yet a fundamental limitation persists across all of these systems:

[†] Corresponding author.



Fig. 1: Examples of geometry artifacts in generated videos and our results.

a lack of explicit geometric supervision during training. This results in emerging geometric artifacts (as shown in Fig. 1): object deformation, spatial drift, flickering, depth violations, physically implausible perspective changes, etc.

While closed-source models [4, 14, 37] tend to exhibit stronger geometric realism than their open-source counterparts, this advantage is largely attributable to training at an unprecedented scale: by fitting massive corpora of natural video, these models allow implicit geometric priors to emerge as a byproduct of sheer data volume. Some recent works have attempted a more principled alternative by incorporating explicit geometric supervision such as conditioning on depth maps [10, 21, 48] or camera poses [17, 44, 45]. However, such approaches are fundamentally constrained by data accessibility: geometric ground truth of this kind is rarely paired with the Internet-scale corpora on which modern video generators are trained, rendering explicit supervision impractical for open-source research and infeasible for customized downstream applications.

In light of these limitations, an increasingly adopted strategy for improving generative diffusion models is to align or tilt them toward an informative reward (either a verifiable reward [7] or a learned reward model [13, 47]). Instead of relying solely on pretraining objectives, an external reward can help capture desired structural or perceptual properties to steer generation. Existing approaches primarily focus on *post-hoc alignment* [31, 33], such as Supervised Fine-Tuning (SFT) or Reinforcement Learning (RL), where model parameters are tilted to reward-favoured distribution. Alternatively, without additional

training, reward-based Test-Time Scaling (TTS) [34, 51] can improve generation by using rewards at inference time to guide search, re-ranking, or iterative refinement among candidate rollouts. Such a paradigm also applies to geometric aspects: recent works [1, 11, 26, 31] have demonstrated that reward-based post-training (e.g., RL) can effectively improve video generation quality, revealing that the success of reward-based alignment hinges on two key components: (a) a well-designed reward function that faithfully captures the target property (e.g., geometric realism), and (b) a procedure that effectively leverages such reward for alignment, either through parameter-updating post-hoc optimization (e.g., SFT/RL) or inference-time optimization (e.g., TTS).

In this work, unlike prior approaches that rely on handcrafted, rule-based rewards (e.g., [11, 12, 26, 31, 33]), we propose a geometry-based reward derived from a pretrained geometric foundation model VGGT [42]. We compute the pointwise reprojection error as a reward signal and additionally leverage the strong representation of VGGT to identify geometrically meaningful regions for guiding sampling. Based on this reward, we curate a diverse prompt suite spanning a wide range of scene types, viewing angles, and camera trajectories, and generate paired videos with contrastive levels of geometric consistency. These pairs form our **GB3DV-25k** dataset, which provides scalable preference data for geometry-aware alignment. We then apply our reward across multiple downstream settings and model families to demonstrate its effectiveness: (a) *post-hoc alignment*: SFT and preference-based RL of a bidirectional diffusion model (using our dataset); and (b) *test-time scaling*: TTS applied to both a bidirectional diffusion model and a causal autoregressive video model, the latter of which, to the best of our knowledge, has not been previously explored. Within the TTS setting for the causal model, we also explore several reward-guided search strategies enabled by the causal rollout mechanism, showing that autoregressive generation supports more structured inference-time exploration, yielding improved scaling behavior.

Our contributions can be summarized as follows:

- We propose a geometry-based reward that measures multi-view consistency via cross-frame pointwise reprojection error using a pretrained geometric foundation model, together with a geometry-aware sampling strategy that focuses evaluation on reliable, semantically meaningful regions.
- Leveraging this reward, we curate **GB3DV-25k**, a preference dataset of 25,600 geometry-ranked video pairs spanning diverse scenes and camera motions, and show that SFT and RL improve the geometric consistency of bidirectional video diffusion models.
- We study reward-guided TTS for video generation, which pioneers TTS for streaming (causal autoregressive) video models. We propose a general TTS framework and instantiate three efficient search strategies, i.e. *Search on Start*, *Search on Path*, and *Beam Search*, that use our reward as a path verifier to improve geometric realism without retraining.

2 Related Works

2.1 Video Diffusion Models

Recent video diffusion models (VDMs) are predominantly built upon Diffusion Transformer (DiT) architectures [36] operating in compressed spatiotemporal latent spaces. They follow a paradigm that first encode video through a 3D variational autoencoder (VAE) that applies joint spatial and temporal compression, then denoise the resulting latent representation with a transformer backbone trained under flow matching [30] or rectified flow objectives [32]. According to the token processing approach, video generation models can be further categorized into *bidirectional* and *causal autoregressive* ones. Bidirectional models [4, 12, 14, 24, 25, 37, 41, 49] attend to all space-time tokens jointly via full 3D spatiotemporal attention [4, 15, 18, 49], yielding high visual fidelity and strong long-range temporal coherence. However, this design requires denoising all frames simultaneously, leading to attention cost that scales quadratically with the total token count and latency proportional to video length. Causal autoregressive models [6, 22, 50, 53] instead generate frames sequentially with KV caching, trading full attention for low-latency streaming. Our test-time scaling strategy is applied on top of causal models in this family, exploiting their sequential generation structure to enable structured inference-time search without any retraining.

2.2 Reward-Tilted Generation

Reward-based methods steer generative models by transferring the reward information into the generative distribution through three paradigms. The first is *gradient-based reward fine-tuning*, where model weights are updated to maximize a reward signal, either via supervised fine-tuning (SFT) on high-reward samples [27] or via direct backpropagation through differentiable rewards [9, 31]. The second is *particle-based scaling methods*, which search over a discrete set of generation candidates, via Best-of- N sampling (BoN) [40] or Test-Time Scaling (TTS) [34, 51], and select the highest-reward output without updating model parameters. The third is *reinforcement learning*, where policy optimization objectives such as RLHF [8, 35], DPO [38], and GRPO [39] are used to align the model with a reward signal. Recent works have explored distilling geometric priors into video diffusion models via RL-based alignment. Epipolar-DPO [26] uses Sampson epipolar distance as a geometry-aware DPO preference signal, while concurrent work VideoGPA [11] extends this to scene-level by reconstructing videos with VGGT [42] and measuring pixel-level discrepancy via re-rendering. Additional related work is discussed in the supplementary material.

3 Preliminaries

3.1 Video Generation via Flow Matching

Modern video diffusion models are increasingly built upon *rectified flow* [32], a framework that defines a straight-line transport between a Gaussian noise

distribution $p_1 = \mathcal{N}(\mathbf{0}, \mathbf{I})$ and the data distribution p_0 . Given a clean video sample $x_0 \sim p_0$ and noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the noisy latent at time $t \in [0, 1]$ is defined by the linear interpolation: $x_t = (1 - t)x_0 + t\epsilon$.

A neural network $v_\theta(x_t, t)$ is trained to predict the velocity field that transports x_t back toward the data manifold by minimizing the flow-matching objective: $\mathcal{L}_{\text{FM}} = \mathbb{E}_{t, x_0, \epsilon} [\|v_\theta(x_t, t) - (x_0 - \epsilon)\|^2]$. At inference time, a clean sample is recovered by numerically integrating the learned velocity field from $t = 1$ to $t = 0$ via an ODE solver. This formulation underlies several state-of-the-art open-source video generators and serves as the backbone for the preference alignment procedure described in Sec. 5.

3.2 Camera Re-Projection

We briefly review the standard pinhole camera model that underpins our geometric reward computation. Each frame i in a video sequence is associated with a camera intrinsic matrix $\mathbf{K}_i \in \mathbb{R}^{3 \times 3}$ and an extrinsic transformation $[\mathbf{R}_i \mid \mathbf{t}_i] \in \mathbb{R}^{3 \times 4}$, where $\mathbf{R}_i$ and $\mathbf{t}_i$ denote the rotation and translation from world coordinates to the camera frame. In our work, these camera parameters and per-frame depth maps are estimated in a feed-forward manner by geometric foundation models [28, 29, 42, 43], which jointly predicts camera intrinsics, extrinsics, and dense depth from a sequence of images without recursive optimization.

Given a 2D pixel location (u_k, v_k) in frame i with associated depth d_k , the corresponding 3D point in world coordinates is obtained via back-projection:

$$\mathbf{X}_k^w = \mathbf{R}_i^{-1} \left(d_k \mathbf{K}_i^{-1} [u_k, v_k, 1]^T - \mathbf{t}_i \right), \quad (1)$$

where $\mathbf{X}_k^w = (X_k, Y_k, Z_k)^T$ denotes the world-space coordinate. To reproject this 3D point into a target frame j with camera parameters $(\mathbf{K}_j, \mathbf{R}_j, \mathbf{t}_j)$, we transform it into the target camera space and apply the perspective projection:

$$\mathbf{X}_k^c = \mathbf{R}_j \mathbf{X}_k^w + \mathbf{t}_j, \quad \left[\hat{u}_k^{(j)}, \hat{v}_k^{(j)}, 1 \right]^T = \frac{1}{Z_k^c} \mathbf{K}_j \mathbf{X}_k^c, \quad (2)$$

where $\mathbf{X}_k^c = (X_k^c, Y_k^c, Z_k^c)^T$ is the camera-space coordinate and Z_k^c is the corresponding depth in frame j . The reprojected pixel $\hat{\mathbf{p}}_k^{(j)} = (\hat{u}_k^{(j)}, \hat{v}_k^{(j)})^T$ serves as a prediction of where point k should appear in frame j , and its deviation from the tracker-predicted correspondence constitutes our reward signal (Sec. 4.2).

4 Geometry-Based Reward

To align video generative models with geometric preference, we introduce Geometry-Based Reward, which takes advantage of an existing geometric foundation model (VGGT [42]) to conduct dense reconstruction and evaluate three-dimensional consistency by point-wise reprojection error. While dense reconstruction provides a geometry-grounded representation of the video, simply computing pixel-space error via warping frame to frame would introduce significant interference.




Fig. 2: Overview of our framework, consisting of two components: (a) **Geometric-Based Reward Model**: a geometry-aware sampling (GAS) module leverages global attention of VGGT to identify salient patches and computes cross-frame pointwise reprojection error; (b) **Geometric Preference Alignment**: the model is aligned via SFT [27] and DPO [31] on a bidirectional model, or test-time scaling (TTS) with our reward as a path verifier on a causal model [53].

To remedy the issue, as shown in Fig 2-(left), we propose a *geometry-aware sampling strategy* that filters out distractors and *computes reprojection error* in a pointwise manner to capture three-dimensional inconsistencies.

4.1 Geometry-Aware Sampling

Previous studies [3, 5, 16, 20] show that the global attention layers of VGGT excavate global information across input frames. For geometric context, we find that the shallow global-attention layers constantly emphasize geometrically meaningful areas (Fig. 3, middle row), as they process more geometric properties of the scene. We leverage this to select geometrically critical areas and filter out irrelevant regions (e.g., sky or ground).

Given a video sequence $\mathcal{I} = \{I_i\}_{i=0}^N$, we pass it through alternating-attention layers [42] and extract the shallow layer to compute the geometric-attention score. Anchoring frame i as the query and the others j as keys, we compute the scaled dot-product attention of the query tokens $\mathbf{Q}_i$ to key tokens $\mathbf{K}_j$:

$$\mathbf{A}_i = \frac{1}{N-1} \sum_{j \neq i}^{N-1} \text{softmax} \left(\frac{\mathbf{Q}_i \mathbf{K}_j^T}{\sqrt{d}} \right), \quad (3)$$

where d is the attention dimension. These token-level scores are averaged across heads and summed over all frames $j \neq i$, then up-sampled to full resolution



Fig. 3: Visualization of Geometry-Aware Sampling, which shows that the global attention of VGGT naturally captures the background geometry. We select top- τ percentage of attention-emphasized patches and sample at the center of each patch.

(H, W) via bilinear interpolation and normalized to $[0, 1]$ to produce a geometric attention heatmap $\mathbf{M}_i$.

We partition each frame into non-overlapping $p \times p$ patches, with each patch’s value being the mean of $\mathbf{M}_i$ within it, and select the top $\tau\%$ by attention value. For each selected patch, we take its center pixel as a sampling point (Fig. 3, bottom row), yielding 2D point locations $\mathcal{P}_k^{(i)} = \{(u_k^{(i)}, v_k^{(i)})\}_{k=1}^{K^{(i)}}$ per frame i , where $K^{(i)}$ is the number of retained patches. By default we set $p = 4$ and $\tau = 20\%$ in all experiments.

4.2 Pointwise Reprojection Error

After identifying geometrically critical regions in each frame, we can establish point correspondences across frames and compute pointwise reprojection errors to quantify geometric consistency.

Point Tracking and Correspondence. For each reference frame i with sampled points $\mathcal{P}_i$, we employ a tracking module to establish correspondences by querying the tracker with points $\mathcal{P}_i$ against all other frames j . This yields tracked positions $\mathbf{p}_k^{(j)} = (u_k^{(j)}, v_k^{(j)})^T$ for each point k across all frames $j \in \{0, \dots, N-1\}$, along with confidence scores $c_k^{(j)} \in [0, 1]$.

Unprojection to 3D Space. For each query point $\mathbf{p}_k$ in reference frame i , we retrieve its depth value $d_k = D_i[v_k, u_k]$ from the VGGT-predicted depth map D_i , together with the corresponding camera parameters $(\mathbf{K}_i, \mathbf{R}_i, \mathbf{t}_i)$. The 3D world coordinate $\mathbf{X}_k^w$ is then recovered via back-projection as defined in Eq. 1.

Reprojection and Error Computation. To validate geometric consistency, we reproject each 3D point $\mathbf{X}_k^w$ into target frame j using Eq. 2, yielding the geometry-grounded estimate $\hat{\mathbf{p}}_k^{(i \rightarrow j)}$. We apply validity filtering of sampled points to ensure robust error computation. The final reprojection error metric is the mean L2 distance over all valid point-frame pairs:

$$\mathcal{E}_{\text{reproj}} = \frac{1}{|\mathcal{V}|} \sum_{(k,i,j) \in \mathcal{V}} \|\hat{\mathbf{p}}_k^{(i \rightarrow j)} - \mathbf{p}_k^{(j)}\|_2, \quad (4)$$

where $\mathcal{V}$ denotes the set of valid point-frame pairs. A lower reprojection error indicates stronger geometric consistency between the tracker’s correspondences and the foundation model’s predicted geometry.

5 Geometry-Guided Preference Alignment

Once a reliable reward model is available, it can be used to influence video generation in several ways: (a) *online optimization*, where the model is updated using reinforcement learning with reward feedback, such as RLHF [8, 35] and GRPO [39]; (b) *offline optimization*, which includes post-hoc preference learning methods such as DPO [38] or supervised fine-tuning [27] on offline generated data; and (c) *test-time optimization*, where the reward guides search, ranking, or iterative refinement of candidate generations without updating model parameters. Due to limitations in model size and computational speed, most video generation methods adopt *offline optimization* [11, 26], particularly SFT or DPO, to incorporate reward knowledge into the generation process.

5.1 Post-hoc Preference Alignment

Preference Data Construction. For offline optimization, a dataset must first be constructed using the reward signal to selectively curate or annotate samples. We construct our **GB3DV-25k** dataset as follows: Given a fixed set of random seeds $\{s_i\}_{i=0}^N$ and condition vector c , we generate a video set via video diffusion $x_0 = G_\theta(z_i, c)$, $z_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ where z_i is the Gaussian noise generated with seed s_i . We then leverage the proposed geometry-based reward as evaluator r_{geo} to assess the videos. To magnify the geometric difference signal across videos, we pick the best and worst samples for each prompt and construct the (x_0^w, x_0^l) pair as our preference data.

Supervised Fine-Tuning. The most direct way to leverage a reward signal is to perform supervised fine-tuning (SFT) in the reward-selected dataset $\mathcal{D}^*$. In this setting, we adopt LoRA [19], a parameter-efficient fine-tuning (PEFT) approach, on the video generator and train the model v_θ using the same diffusion/flow matching objective:

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{t, x_0 \in \mathcal{D}^*, \epsilon} \left[\|v_\theta(x_t, t) - (x_0 - \epsilon)\|^2 \right]. \quad (5)$$

DPO with Geometric Preference. Direct Preference Optimization (DPO) is an offline reinforcement learning method that aligns the policy using pairwise preferences under the Bradley–Terry model. It leverages preference comparisons to directly optimize the policy without requiring a globally normalized reward function. Given preference pairs (x_0^w, x_0^l) , the geometric preference DPO optimizes the parameters θ via:

$$\max_{\theta} \mathbb{E}_{(x_0^w, x_0^l) \sim \mathcal{D}} \left[\log \sigma \left(\beta \log \frac{p_{\theta}(x_0^w | c)}{p_{\text{ref}}(x_0^w | c)} - \beta \log \frac{p_{\theta}(x_0^l | c)}{p_{\text{ref}}(x_0^l | c)} \right) \right], \quad (6)$$

where β is the coefficient that controls the deviation from the reference policy p_{ref} , and $\mathcal{D}$ denotes the preference dataset containing $x_0^w \succ x_0^l$ ranked by r_{geo} .

In our work, we adopt the formulation of Flow-DPO [31], which reformulates the objective for rectified flow models:

$$\mathcal{L} = -\mathbb{E} \left[\log \text{sigmoid} \left(-\frac{\beta_t}{2} \left(\|v^w - v_{\theta}(\mathbf{x}_t^w, t)\|^2 - \|v^w - v_{\text{ref}}(\mathbf{x}_t^w, t)\|^2 - \left(\|v^l - v_{\theta}(\mathbf{x}_t^l, t)\|^2 - \|v^l - v_{\text{ref}}(\mathbf{x}_t^l, t)\|^2 \right) \right) \right) \right], \quad (7)$$

where v_{ref} and v_{θ} are the velocities predicted by the model, $\beta_t = \beta(1 - t^2)$ is the training weight that depends on the noise level, and $x_t^* = (1 - t)x_0^* + t\epsilon$ is the noisy latent of the preference pair.

To prevent mode collapse, we introduce two auxiliary loss terms: a first-order term that penalizes static motion and a second-order term that encourages overall smoothness, applied with opposite signs:

$$\mathcal{L}_{\text{aux}} = -\mathbb{E}[\text{Var}_t(\hat{x}_0)] + \gamma \cdot \mathbb{E}[\|\Delta_t^2 \hat{x}_0\|^2], \quad (8)$$

where $\hat{x}_0$ is the reconstructed clean sample, $\Delta_t^2 \hat{x}_0$ denotes the second-order temporal difference of $\hat{x}_0$, and $\gamma > 0$ is the weighting coefficient. The final loss of DPO is: $\mathcal{L}_{\text{total}} = \mathcal{L} + \lambda \mathcal{L}_{\text{aux}}$, where λ controls the strength of the penalty.

5.2 TTS with Geometric Preference

In contrast to most *post-hoc* alignment methods that require constructing a dataset beforehand, Test-Time Scaling (TTS) optimizes the reward signal directly at inference time. It aims to identify an optimal sample using a verifier $\mathcal{V}$ together with a search procedure f :

$$f : \mathcal{V} \times G_{\theta} \times \{\mathbb{R}^{N \times H \times W \times C} \times \mathbb{R}^d\}^N \rightarrow \mathbb{R}^{N \times H \times W \times C}, \quad (9)$$

where G_{θ} is the pretrained generator and $\mathcal{V}$ evaluates the quality of candidates.

For bidirectional models, which denoise all spatial-temporal tokens jointly, TTS reduces to a Best-of- N protocol similar to most image generation TTS scenarios [34, 51]: N candidates are generated in parallel, and the highest-rewarded sample is returned. Causal autoregressive models, by contrast, generate frames

sequentially, where each frame x_t is conditioned on the previously generated prefix $x_{<t}$. This temporal Markov structure exposes a richer search space: one can intervene at each step to prune low-quality paths and explore diverse trajectories, going beyond the simple Best-of- N selection that only retains a narrow search space.

To exploit this richer search space, we instantiate TTS on a causal autoregressive model [53] and reformulate the problem as searching for a generation path in a space spanned by the seed axis and temporal axis, with the goal of finding the highest-rewarded sample. Applying brute-force search guarantees an optimal solution, but it yields intractable $\mathcal{O}(S^N)$ complexity for a video clip of N frames with a seed search range of size S . To address this, we propose three search algorithms (shown in Fig. 2, right bottom) that all achieve a reasonable complexity but with different dynamics:

Search on Start (SoS). The algorithm operates along the seed axis. Given a fixed set of S candidate seeds $\{s_1, s_2, \dots, s_S\}$, the algorithm performs a complete forward pass for each seed independently and evaluates the resulting video clip using the reward model r_{geo} . The seed yielding the highest reward is selected, and its corresponding output clip is returned:

$$s_{best} = \arg \max_{s_i} r_{geo}[G_\theta(z_i, c)], \quad i = 1, \dots, S. \quad (10)$$

Search on Path (SoP). The algorithm proceeds along the temporal axis, dynamically selecting from a set of seed candidates S at each time step t . When generating the next frame, it iterates over S and selects the optimal seed s_{best}^t that yields the highest reward. The reward model evaluates frames within a sliding context window $\mathcal{W}$ spanning the preceding w frames, including the current one. The seed path that achieves the highest cumulative reward is then returned:

$$s_{best}^t = \arg \max_{s_i^t} r_{geo}[\mathcal{W} \cup \{G_\theta(z_i^t, c)\}], \quad i = 1, \dots, S. \quad (11)$$

Beam Search (BS). The algorithm generalizes both SoS and SoP by maintaining K candidate paths throughout generation. At each time step, it evaluates $K \times S$ child nodes spawned from all current paths and retains only the top- K nodes, pruning the rest. The path achieving the highest cumulative reward is returned. The objective follows the same formulation as Eq. 11, with the distinction that K optimal seeds s_{best}^t are retained at every step rather than one.

It's noteworthy that the previous two search plans can be viewed as special configurations of BS: SoS corresponds to $(K, 1)$ and SoP corresponds to $(1, S)$, with complexities of $\mathcal{O}(KN)$ and $\mathcal{O}(SN)$. The beam search algorithm supports flexible (K, S) configurations, yielding a general complexity of $\mathcal{O}(KSN)$.

6 Experiments

6.1 Implementation Details

Dataset Curation. Preference alignment requires training pairs with notable differences in geometric consistency. To ensure sufficient diversity, we generate 10 samples per prompt using bidirectional CausVid [50]. To construct the prompt suite, we source scene references from two representative datasets: RealEstate10k [52] for indoor scenes and GLDv2 [46] for outdoor scenes. We employ Qwen3-VL [2] to generate detailed scene captions, explicitly instructing the model to describe both static and dynamic objects, diverse camera movements, and varying shooting angles. The resulting prompt suite comprises 2,560 entries, from which CausVid generates a total of 25,600 video clips, constituting our curated GB3DV-25k dataset.

Evaluation Metrics. For geometric consistency, we adopt two complementary protocols: 3D reconstruction quality for dense evaluation and multi-view consistency metrics for sparse evaluation. For 3D reconstruction quality, we apply VGGT [42] to uniformly sampled frames to obtain depth maps and camera poses, reproject the recovered 3D point cloud into each target frame, and measure re-projection fidelity via PSNR, SSIM, and LPIPS. For multi-view consistency, we compute three scores on the evaluated videos: epipolar consistency (EPI), reprojection-pixelwise (RPX), and reprojection-pointwise (RPT). For overall video quality, we use VBench [23], a comprehensive benchmark covering sub-ject consistency (SC), background consistency (BC), motion smoothness (MS), dynamic degree (DD), aesthetic quality (AQ), and imaging quality (IQ). All VBench scores reported in our tables are normalized using the empirical minimum and maximum values specified in the VBench paper [23].

Baselines. We compare our methods against the following baselines:

Base Model. We use the Causvid [50] for TTS of bidirectional experiments (Sec. 6.2), Causal-Forcing [53] for TTS of streaming video experiments (Sec. 6.3), and Wan2.1-T2V-1.3B [41] for post-hoc alignment experiments (Sec. 6.4).

Epipolar. Built directly on the official code of [26], assessing geometric consistency via epipolar constraints quantified by the Sampson distance.

Reproj-Pix. A faithful reproduction of VideoGPA [11], which employs VGGT to estimate camera geometry and quantifies the reward via pixelwise intensity differences between the source and warped frames.

6.2 TTS of Bidirectional Video Generation

Evaluation Setups. We evaluate the effectiveness of test-time scaling on bidirectional video generation by adopting the best-of- N sampling protocol. During evaluation, each reward model independently selects the highest-scoring video from N generated candidates. For generality, the experiments are conducted on the full-scale GB3DV-25k dataset and N is set to 10.

Table 1: Evaluation of Test-Time Scaling via Best-of-N Sampling.

Method	Geometric Consistency Evaluation						Overall Video Quality Evaluation (VBench %)					
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	EPI $\downarrow$	RPX $\downarrow$	RPT $\downarrow$	SC $\uparrow$	BC $\uparrow$	MS $\uparrow$	AQ $\uparrow$	IQ $\uparrow$	Total $\uparrow$
Baseline [50]	19.68	0.6381	0.3604	5.553	0.9738	4.706	94.16	92.36	97.16	57.28	75.68	83.33
Epipolar [26]	<u>22.45</u>	<u>0.7557</u>	<u>0.2432</u>	-	<u>0.9546</u>	2.815	<u>96.22</u>	94.02	98.14	<u>58.20</u>	75.91	<u>84.50</u>
Reproj-Pix [11]	21.07	0.7267	0.3036	<u>4.549</u>	-	<u>3.473</u>	95.21	92.97	97.48	58.04	76.14	83.97
Reproj-Pts (Ours)	22.66	0.7665	0.2330	3.442	0.9539	-	96.39	<u>93.95</u>	<u>97.96</u>	58.23	<u>76.09</u>	84.52

Evaluation Results. Quantitative results are shown in Tab. 1. We omit each method’s score on its own selection metric (marked "-"), as such self-referential entries are trivially optimal and obscure a fair comparison. Our *Reproj-Pts* reward achieves the best results on all three 3D reconstruction metrics and the best EPI and RPX among multi-view metrics, showing that pointwise reprojection error is a more robust and physically grounded signal than pixel-space alternatives. The consistent gain over Reproj-Pix [11] confirms the benefit of decoupling geometric error from pixel intensity. For overall quality, our method attains the highest total VBench score (84.52%) while leading on subject consistency and aesthetic quality, indicating that optimizing geometric consistency does not sacrifice perceptual quality.

6.3 TTS of Streaming Video Generation

Evaluation Setups. We apply test-time scaling to streaming video generation and evaluate the three proposed search algorithms: Search on Start (SoS), Search on Path (SoP), and Beam Search (BS). We conduct a budget scaling study spanning budgets from 1 to 16 (i.e. budget of 4 means 4 seeds for SoS and SoP, 2 child nodes and top-2 for BS) on a 16-clip subset of GB3DV-25k to investigate the scaling behavior of different searching schemes of TTS.

Evaluation Results. Budget scaling results are shown in Fig. 4. PSNR, SSIM, LPIPS, and RPT are for 3D consistency metrics, and SC, BC, MS, AQ, and IQ are the perceptual quality dimensions of VBench. We observe two findings. (1) **All variants exhibit scaling behavior:** as the budget increases, all three methods improve consistently on both geometric and perceptual metrics, confirming that our geometry-based reward offers a meaningful search signal. (2) **Budget configuration shapes the type of improvement:** SoS allocates $K > 1$ parallel paths for greater diversity but lacks temporal refinement, achieving the lowest RPT yet weaker results elsewhere; SoP maintains $S > 1$ seeds per step for stable fine-grained gains, yielding the best VBench trend; Beam Search combines both ($K > 1, S > 1$) and attains the strongest 3D reconstruction. Qualitatively (Fig. 5), the baseline shows geometric artifacts and incorrect perspective in later frames, while all three strategies produce geometrically coherent videos throughout. Note that SoP shares the baseline’s initial seed and thus an identical first frame, whereas SoS and Beam Search optimize the seed before generation and explore a larger space from the first frame.




Fig. 4: Budget Evaluation of TTS. All three methods—SoS, SoP, and Beam Search—demonstrate scaling tendencies as the search budget increases. Beam Search prevails in 3D metrics, while SoP achieves the best overall performance.

6.4 Post-hoc Alignment of Bidirectional Video Generation

Evaluation Setups. We apply post-hoc alignment to a bidirectional DiT [41] using our geometry-based reward, evaluating two complementary strategies: supervised fine-tuning (SFT) and Direct Preference Optimization (DPO). For both, we adopt Low-Rank Adaptation (LoRA) [19] with rank $r = 64$ and $\alpha = 128$ applied to the q, k, v, and o projection modules of the DiT. For the auxiliary loss terms, we configure $\lambda = 0.1$ and $\gamma = 0.01$.

Results of post-hoc alignment. Quantitative results are shown in Tab. 2. SFT already yields notable improvements over the baseline across most geometric metrics, as fine-tuning on high-reward samples tilts the generation distribution toward geometrically preferred modes. DPO further improves upon SFT by explicitly contrasting winning and losing samples, achieving stronger geometric consistency across 3D reconstruction and multi-view metrics. Among DPO variants, our Reproj-Pts reward attains the best SSIM (0.7977), LPIPS (0.1789), and EPI (2.127), outperforming the Epipolar baseline on the metrics most reflective of dense geometric coherence. For overall video quality, both SFT and DPO consistently improve over the baseline on subject and background consistency, while our DPO (Reproj-Pts) achieves the best SC (97.05%), BC (95.25%), and IQ (76.63%), demonstrating that geometric preference alignment enhances perceptual quality alongside geometric fidelity.



Fig. 5: Qualitative Results of Test-Time Scaling. The baseline exhibits geometric artifacts (highlighted by $\times$), wrong perspective relation as shown in the last two frames. Our approach, whether optimizing over the initial seed (SoS), selecting frame-by-frame along the temporal axis (SoP), or applying beam search (BS), consistently produces geometrically coherent videos with no visible artifacts (highlighted by $\checkmark$).

Table 2: Evaluation of Post-hoc Alignment.

Method	Geometric Consistency Evaluation						Overall Video Quality Evaluation (VBench %)					
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	EPI $\downarrow$	RPIX $\downarrow$	RPT $\downarrow$	SC $\uparrow$	BC $\uparrow$	MS $\uparrow$	DD $\uparrow$	AQ $\uparrow$	IQ $\uparrow$
Baseline [41]	22.45	0.7548	0.2243	2.832	0.998	1.783	95.98	94.43	97.61	27.34	<u>57.74</u>	76.30
+ SFT (Reproj-Pts)	23.52	0.7927	0.1842	2.337	<u>1.003</u>	2.257	96.97	95.15	97.96	15.62	57.73	<u>76.58</u>
+ SFT + DPO (Epipolar)	23.57	<u>0.7973</u>	<u>0.1818</u>	<u>2.187</u>	1.018	1.385	<u>96.98</u>	<u>95.16</u>	<u>97.88</u>	25.39	57.76	76.52
+ SFT + DPO (Reproj-Pts)	<u>23.54</u>	0.7977	0.1789	2.127	1.022	<u>1.424</u>	97.05	95.25	97.77	<u>25.78</u>	57.71	76.63

6.5 Ablation Studies

Regularization in Post-hoc Alignment. The geometry reward inherently favors static frames, leading to dynamic degree (DD) collapse under DPO. The auxiliary loss in Eq. 8, with static penalty λ and smoothness weight γ , counteracts this collapse. Tab. 3 reports deltas w.r.t. $\lambda = 0, \gamma = 0$, isolating the DD recovery attributable to Eq. 8 from the DPO objective itself. A moderate λ restores DD with negligible quality cost, whereas an overly large γ harms motion smoothness.

Table 3: Ablation of Regularization Terms in Eq. 8.

λ	SC $\uparrow$	BC $\uparrow$	MS $\uparrow$	DD $\uparrow$	AQ $\uparrow$	IQ $\uparrow$
0.1	+0.02	-0.04	+0.00	+0.78	-0.02	+0.02
0.5	+0.05	-0.08	+0.02	+0.78	-0.06	-0.01
1.0	-0.00	-0.02	+0.00	+0.39	-0.15	-0.00
10.0	-0.02	+0.17	+0.08	+1.56	-0.01	-0.18

γ	SC $\uparrow$	BC $\uparrow$	MS $\uparrow$	DD $\uparrow$	AQ $\uparrow$	IQ $\uparrow$
0.1	-0.03	+0.04	-0.01	-0.39	+0.01	+0.04
0.5	-0.04	+0.05	-0.02	+0.00	+0.04	+0.05
1.0	+0.01	+0.03	-0.03	+0.00	-0.03	+0.02
10.0	-0.05	-0.10	-0.12	+0.00	-0.12	+0.02

Table 4: Ablation of Geometry-Aware Sampling (GAS).

Setting	SC $\uparrow$	BC $\uparrow$	MS $\uparrow$	DD $\uparrow$	AQ $\uparrow$	IQ $\uparrow$
Uniform	-0.06	-0.00	+0.08	-4.30	-0.11	+0.02
Deep Attn	-0.12	+0.17	+0.05	-2.34	-0.26	-0.19
$p=8$	-0.02	-0.02	+0.01	-1.17	-0.02	+0.02
$\tau=5\%$	-0.18	-0.32	-0.06	+1.17	+0.25	+0.11
$\tau=40\%$	-0.08	-0.03	-0.01	-0.78	+0.03	-0.07

Setting	SC $\uparrow$	BC $\uparrow$	MS $\uparrow$	DD $\uparrow$	AQ $\uparrow$	IQ $\uparrow$
Saliency	-0.17	+0.04	+0.08	-4.69	+0.02	-0.16
$p=2$	-0.27	-0.37	-0.06	+2.34	+0.30	+0.11
$p=16$	-0.07	+0.03	-0.03	-0.39	+0.15	-0.04
$\tau=10\%$	-0.22	-0.26	-0.04	+1.56	+0.19	+0.08
$\tau=80\%$	-0.08	-0.03	-0.01	-0.78	+0.03	-0.07

Geometry-Aware Sampling (GAS). Tab. 4 ablates the sampling strategy, patch size p , and ratio τ , with deltas relative to our default (shallow attention, $p = 4$, $\tau = 20\%$). Attention-based selection clearly surpasses uniform and saliency sampling, and the default configuration of p and τ best balances geometric coverage against perceptual quality.

7 Conclusions

In this paper, we introduced VIGOR, a Video Geometry-Oriented Reward framework that addresses the lack of explicit geometric supervision in current video diffusion models. Leveraging pretrained geometric foundation models, we proposed a physically grounded reward based on pointwise cross-frame reprojection error, coupled with a geometry-aware sampling strategy for robustness in low-texture and non-semantic regions. Extensive experiments show that VIGOR mitigates common temporal artifacts such as object deformation, spatial drift, and depth violations. We further validated its versatility across post-training alignment and inference-time optimization for both bidirectional and causal architectures, providing a practical and scalable solution for geometrically consistent video generation.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant No. 62302240), the Beijing Major Science and Technology Project (Grant No. Z251100007125021), and by Hi! PARIS through the Hi! PARIS Chair 2024 held at École Polytechnique (LIX). Additional computational resources were provided by the Supercomputing Center of Nankai University and the IDRIS High-Performance Computing facilities (under allocation 2026-AD011014300R3, courtesy of GENCI).

References

1. Asim, M., Wewer, C., Wimmer, T., Schiele, B., Lenssen, J.E.: MET3R: Measuring multi-view consistency in generated images. In: IEEE Conf. Comput. Vis. Pattern Recog. (CVPR). pp. 6034–6044 (2025)
2. Bai, S., Cai, Y., Chen, R., Chen, K., Chen, X., et al.: Qwen3-vl technical report. Tech. rep., Alibaba Group (2025), arXiv:2511.21631
3. Bratulić, J., Mittal, S., Brox, T., Rupperecht, C.: On geometric understanding and learned data priors in vggd (2025), arXiv:2512.11508
4. Brooks, T., Peebles, B., Holmes, C., DePue, W., et al.: Video generation models as world simulators. Tech. rep., OpenAI (2024), <https://openai.com/research/video-generation-models-as-world-simulators>, accessed: 2026-06-26
5. Cao, Y., Wu, F., Chen, D.Z., Zhong, Y., Hong, L., Xu, D.: Vggt-det: Mining vggt internal priors for sensor-geometry-free multi-view indoor 3d object detection. In: IEEE Conf. Comput. Vis. Pattern Recog. (CVPR) (2026)
6. Chen, B., Monso, D.M., Du, Y., Simchowitz, M., Tedrake, R., Sitzmann, V.: Diffusion forcing: Next-token prediction meets full-sequence diffusion. In: Adv. Neural Inform. Process. Syst. (NeurIPS) (2024)
7. Chen, J., Huang, Y., Lv, T., Cui, L., Chen, Q., Wei, F.: Textdiffuser: Diffusion models as text painters. In: Adv. Neural Inform. Process. Syst. (NeurIPS) (2023)
8. Christiano, P., Leike, J., Brown, T.B., Martic, M., Legg, S., Amodei, D.: Deep reinforcement learning from human preferences. In: Adv. Neural Inform. Process. Syst. (NeurIPS) (2017)
9. Clark, K., Vicol, P., Swersky, K., Fleet, D.J.: Directly fine-tuning diffusion models on differentiable rewards. In: Int. Conf. Learn. Represent. (ICLR). vol. 2024, pp. 4793–4822 (2024)
10. Dai, Y., Jiang, F., Wang, C., Xu, M., Qi, Y.: Fantasyworld: Geometry-consistent world modeling via unified video and 3d prediction (2025), arXiv:2509.21657
11. Du, H., Ye, J., Cong, X., Li, R., Ni, J., et al.: Videogpa: Distilling geometry priors for 3d-consistent video generation (2026), arXiv:2601.23286
12. Gao, Y., Guo, H., Hoang, T., Huang, W., Jiang, L., et al.: Seedance 1.0: Exploring the boundaries of video generation models. Tech. rep., ByteDance (2025), arXiv:2506.09113
13. Ghosh, D., Hajishirzi, H., Schmidt, L.: Geneval: An object-focused framework for evaluating text-to-image alignment. In: Adv. Neural Inform. Process. Syst. (NeurIPS) (2023)
14. Google DeepMind: Veo 3 technical report. Tech. rep., Google DeepMind (2025), <https://storage.googleapis.com/deepmind-media/veo/Veo-3-Tech-Report.pdf>, accessed: 2026-06-26
15. Gupta, A., Yu, L., Sohn, K., Gu, X., Hahn, M., Li, F.F., Essa, I., Jiang, L., Lezama, J.: Photorealistic video generation with diffusion models. In: Eur. Conf. Comput. Vis. (ECCV) (2023)
16. Han, J., Hong, S., Jung, J., Jang, W., An, H., et al.: Emergent outlier view rejection in visual geometry grounded transformers. In: IEEE Conf. Comput. Vis. Pattern Recog. (CVPR). pp. 427–437 (2026)
17. He, H., et al.: Cameractrl ii: Dynamic scene exploration via camera-controlled video diffusion models. In: IEEE Int. Conf. Comput. Vis. (ICCV). pp. 13416–13426 (2025)
18. Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. In: Adv. Neural Inform. Process. Syst. (NeurIPS) (2022)

19. 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. In: *Int. Conf. Learn. Represent. (ICLR)* (2021)
20. Hu, Y., Cheng, C., Yu, S., Guo, X., Wang, H.: VGGT4D: Mining motion cues in visual geometry transformers for 4d scene reconstruction (2025), arXiv:2511.19971
21. Huang, T., Zheng, W., Wang, T., Liu, Y., Wang, Z., Wu, J., et al.: Voyager: Long-range and world-consistent video diffusion for explorable 3d scene generation. *ACM Trans. Graph. (TOG)* **44**, 1 – 15 (2025)
22. Huang, X., Li, Z., He, G., Zhou, M., Shechtman, E.: Self forcing: Bridging the train-test gap in autoregressive video diffusion (2025), arXiv:2506.08009
23. Huang, Z., He, Y., Yu, J., Zhang, F., Si, C., et al.: VBench: Comprehensive benchmark suite for video generative models. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 21807–21818 (2024)
24. Kling Team, Chen, J., Ci, Y., Du, X., Feng, Z., et al.: Kling-omni technical report. Tech. rep., Kuaishou (2025), arXiv:2512.16776
25. Kong, W., Tian, Q., Zhang, Z., Min, R., Dai, Z., et al.: Hunyuanvideo: A systematic framework for large video generative models (2025), arXiv:2412.03603
26. Kupyn, O., Manhardt, F., Tombari, F., Rupperecht, C.: Epipolar geometry improves video generation models (2025), arXiv:2510.21615
27. Lee, K., Liu, H., Ryu, M., Watkins, O., Du, Y., Boutilier, C., Abbeel, P., Ghavamzadeh, M., Gu, S.S.: Aligning text-to-image models using human feedback (2023), arXiv:2302.12192
28. Leroy, V., Cabon, Y., Revaud, J.: Grounding image matching in 3d with MAST3R. In: *Eur. Conf. Comput. Vis. (ECCV)* (2024)
29. Lin, H., Chen, S., Liew, J., Chen, D.Y., Li, Z., Shi, G., Feng, J., Kang, B.: Depth anything 3: Recovering the visual space from any views (2025), arXiv:2511.10647
30. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: *Int. Conf. Learn. Represent. (ICLR)* (2023)
31. Liu, J., Liu, G., Liang, J., Yuan, Z., Liu, X., et al.: Improving video generation with human feedback (2025), arXiv:2501.13918
32. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. In: *Int. Conf. Learn. Represent. (ICLR)* (2022)
33. Lu, Y., Zeng, Y., Li, H., Ouyang, H., Wang, Q., et al.: Reward forcing: Efficient streaming video generation with rewarded distribution matching distillation (2025), arXiv:2512.04678
34. Ma, N., Tong, S., Jia, H., Hu, H., Su, Y.C., et al.: Inference-time scaling for diffusion models beyond scaling denoising steps (2025), arXiv:2501.09732
35. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C.L., et al.: Training language models to follow instructions with human feedback. In: *Adv. Neural Inform. Process. Syst. (NeurIPS)* (2022)
36. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: *IEEE Int. Conf. Comput. Vis. (ICCV)*. pp. 4172–4182 (2023)
37. Polyak, A., Zohar, A., Brown, A., Tjandra, A., Sinha, A., et al.: Movie gen: A cast of media foundation models (2025), arXiv:2410.13720
38. Rafailov, R., Sharma, A., Mitchell, E., Ermon, S., Manning, C.D., Finn, C.: Direct preference optimization: Your language model is secretly a reward model. In: *Adv. Neural Inform. Process. Syst. (NeurIPS)* (2024)
39. Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., et al.: DeepSeekMath: Pushing the limits of mathematical reasoning in open language models (2024), arXiv:2402.03300

40. Verdun, C.M., Oesterling, A.X., Lakkaraju, H., du Pin Calmon, F.: Soft best-of-n sampling for model alignment. *IEEE Int. Symp. Inform. Theory. (ISIT)* pp. 1–6 (2025)
41. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., et al.: Wan: Open and advanced large-scale video generative models (2025), [arXiv:2503.20314](https://arxiv.org/abs/2503.20314)
42. Wang, J., Chen, M., Karaev, N., Vedaldi, A., Rupprecht, C., Novotny, D.: Vggg: Visual geometry grounded transformer. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 5294–5306 (2025)
43. Wang, S., Leroy, V., Cabon, Y., Chidlovskii, B., Revaud, J.: DUS3R: Geometric 3d vision made easy. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 20697–20709 (2024)
44. Wang, X., Courant, R., Christie, M., Kalogeiton, V.: Akira: Augmentation kit on rays for optical video generation. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 2609–2619 (2025)
45. Wang, Z., Yuan, Z., Wang, X., Chen, T., Xia, M., et al.: MotionCtrl: A unified and flexible motion controller for video generation. In: *ACM SIGGRAPH (SIGGRAPH)* (2024)
46. Weyand, T., Araujo, A., Cao, B., Sim, J.: Google Landmarks Dataset v2 – a large-scale benchmark for instance-level recognition and retrieval. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 2572–2581 (2020)
47. Wu, X., Sun, K., Zhu, F., Zhao, R., Li, H.: Human preference score: Better aligning text-to-image models with human preference. In: *IEEE Int. Conf. Comput. Vis. (ICCV)*. pp. 2096–2105 (2023)
48. Xing, J., Xia, M., Liu, Y., Zhang, Y., Zhang, Y., He, Y.Y., Liu, H., Chen, H., Cun, X., Wang, X., Shan, Y., Wong, T.T.: Make-your-video: Customized video generation using textual and structural guidance. *IEEE Trans. Vis. Comput. Graph. (TVCG)* **31**, 1526–1541 (2023)
49. Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., et al.: CogVideoX: Text-to-video diffusion models with an expert transformer. In: *Int. Conf. Learn. Represent. (ICLR)*. vol. 2025, pp. 83048–83077 (2025)
50. Yin, T., Zhang, Q., Zhang, R., Freeman, W.T., Durand, F., et al.: From slow bidirectional to fast autoregressive video diffusion models. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 22963–22974 (2025)
51. Zhang, X., Lin, H., Ye, H., Zou, J., Ma, J., et al.: Inference-time scaling of diffusion models through classical search (2025), [arXiv:2505.23614](https://arxiv.org/abs/2505.23614)
52. Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: Learning view synthesis using multiplane images. *ACM Trans. Graph. (TOG)* **37**(4), 65:1–65:12 (2018)
53. Zhu, H., Zhao, M., He, G., Su, H., Li, C., Zhu, J.: Causal forcing: Autoregressive diffusion distillation done right for high-quality real-time interactive video generation. In: *Int. Conf. on Mach. Learn. (ICML)* (2026)