

AR-CoPO: Align Autoregressive Video Generation with Contrastive Policy Optimization

Dailan He^{1,2} Guanlin Feng² Xingtong Ge^{2,3} Yi Zhang²
Bingqi Ma² Guanglu Song² Yu Liu² Hongsheng Li^{1,4}

¹ CUHK MMLab, China hedailan@link.cuhk.edu.hk

² Vivix Group Limited, China

³ HKUST, China

⁴ Shenzhen Loop Area Institute, China

Abstract. Streaming autoregressive (AR) video generators combined with few-step distillation achieve low-latency, high-quality synthesis, yet remain difficult to align via reinforcement learning from human feedback (RLHF). Existing SDE-based GRPO methods face challenges in this setting: few-step ODEs and consistency model samplers deviate from standard flow-matching ODEs, and their short, low-stochasticity trajectories are highly sensitive to initialization noise, rendering intermediate SDE exploration ineffective. We propose **AR-CoPO** (**A**uto**R**egressive **C**ontrastive **P**olicy **O**ptimization), a framework that adapts the Neighbor GRPO contrastive perspective to streaming AR generation. AR-CoPO introduces chunk-level alignment via a forking mechanism that constructs neighborhood candidates at a randomly selected chunk, assigns sequence-level rewards, and performs localized GRPO updates. We further propose a semi-on-policy training strategy that complements on-policy exploration with exploitation over a replay buffer of reference roll-outs, improving generation quality across domains. Experiments on Self-Forcing demonstrate that AR-CoPO improves both out-of-domain generalization and in-domain human preference alignment over the baseline, providing evidence of genuine alignment rather than reward hacking.

Keywords: Autoregressive Video Generation · RLHF

1 Introduction

Diffusion and flow-matching models [4, 10, 16] have achieved remarkable progress in image and video synthesis [6, 7, 16, 24, 31, 32], delivering high-quality and diverse results. Despite their success, the inference cost of bidirectional generation typically scales linearly with both the number of sampling steps and the target video length. This scaling behavior makes it difficult to deploy these models in low-latency, variable-length, and streaming-generation scenarios. To address this, a growing line of work [2, 5, 27, 30, 33] distills powerful bidirectional video models into causal generators that operate in an autoregressive (AR) and




Fig. 1: AR-CoPO is a reinforcement learning for human preference (RLHF) method, aligning few-step autoregressive (AR) video generative models to better sample quality.




Fig. 2: Left: Training curves comparing SDE-based GRPO and AR-CoPO on Self-Forcing. SDE-based GRPO fails to improve the reward, while AR-CoPO consistently achieves higher scores throughout training. **Right:** Perturbing only the intermediate CM solver noise (Rows 3–5) produces nearly identical outputs, whereas replacing the initial noise (Row 2) causes significant variation, confirming that few-step AR models (*e.g.* Self-Forcing [5]) are near-deterministic and driven primarily by initial noise.

chunk-wise manner. In addition, techniques like distribution matching distillation (DMD) [28, 29] compress the sampling process into just a few steps. These are often implemented via few-step ODE solvers [13, 20] or consistency models, while engineering optimizations such as KV caching further improve throughput and responsiveness.

However, the combination of a streaming AR structure and few-step distillation introduces significant challenges for post-training alignment, particularly for reinforcement learning from human feedback (RLHF) [1, 14, 21]. RLHF is essential for controllable, high-quality video generation. Similar to language modeling, video generators are often aligned to reward models or verifiers that capture instruction following, subject consistency, motion plausibility, and aesthetics. For flow-matching style generators, policy-gradient post-training (*e.g.*, GRPO-like objectives [18, 19]) provides a natural route by viewing the sampling process as policy rollouts and optimizing the induced distribution using reward feedback.

Many practical GRPO variants designed for flow-matching models, however, hinge on an important implementation choice: converting deterministic ODE sampling into a stochastic SDE form to introduce a Markov decision process (MDP). Applying these SDE-based GRPO methods [8, 9, 11, 23, 26] to few-step streaming generators presents significant challenges. First, few-step generators—often formulated as distilled few-step ODEs or consistency models—deviate from standard flow-matching ODEs, making them inherently difficult to train using SDE-based methods designed for continuous flow matching. Second, because these models follow short sampling trajectories with limited stochasticity, they are highly sensitive to initialization noise and model approximation errors. In contrast, SDE-based approaches heavily rely on intermediate noise injections to guide exploration while typically keeping the initial noise frozen, leading to a fundamental mismatch. This paper, therefore, addresses a critical question: How can we effectively align streaming AR video generators using GRPO-like objectives?

We draw inspiration from the recently proposed Neighbor GRPO [3] perspective, which offers a novel reinterpretation: the SDE-GRPO update can be mathematically reformulated as a distance-driven contrastive objective over a set of neighbor candidate trajectories. Crucially, this view suggests that one can obtain informative preference signals without relying on stochastic SDE exploration during sampling. Instead, exploration can be entirely controlled during training by constructing neighborhood candidates around the initial noise of an ODE sampler and defining a softmax distance-based surrogate transition distribution over these candidates.

Based on this insight, we propose an **AutoRegressive Contrastive Policy Optimization (AR-CoPO)** framework for the post-training alignment of streaming AR video generators. To match the structural properties of streaming AR generation, we introduce **chunk-level alignment** as the fundamental unit of neighborhood construction and optimization. This establishes a chunk-level action space that responds to sequence-level rewards, enabling natural and localized credit assignment. Furthermore, in addition to standard on-policy training that encourages exploration, we adopt a **semi-on-policy** training strategy under the AR-CoPO objective to enhance exploitation. By combining these two training paradigms, we effectively improve both out-of-domain performance and overall generation quality. In summary, our contributions are:

1. We propose the AR-CoPO framework, which enables the RLHF of streaming video generators using a contrastive objective over neighborhood candidates.
2. We introduce a chunk-level action space that allows for natural credit assignment from sequence-level rewards, producing localized and controllable reward signals.
3. We propose a semi-on-policy training strategy under the AR-CoPO objective to enhance exploitation. Combined with standard on-policy exploration, this effectively improves both out-of-domain performance and overall generation quality.

2 Related Works

2.1 Autoregressive Video Generation

Flow matching (FM) [10] and related probability flow ODEs provide a deterministic alternative to stochastic diffusion sampling, enabling fast generation via few-step ODE solvers when combined with distillation techniques like distribution matching distillation (DMD) [28, 29]. While ODE determinism enables efficient and stable inference, it limits the stochastic exploration typically required for RL-based alignment. Our work targets alignment mechanisms that enable controllable exploration during training while preserving fast deterministic inference.

To support streaming applications and reduce latency, recent methods distill bidirectional video generators into *causal* autoregressive (AR) generators that synthesize videos chunk by chunk [5, 27, 30, 33]. However, the AR setting introduces new challenges, such as cross-chunk error accumulation and exposure bias due to the mismatch between training and inference contexts. These challenges motivate our focus on developing stable and localized alignment methods tailored for streaming AR generation and short-horizon few-step ODE inference.

2.2 Post-training Flow Alignment and Neighbor GRPO

Post-training alignment for generative models typically leverages human preferences or computable rewards to improve instruction following, aesthetics, and temporal consistency. For diffusion and flow-based generators, reinforcement-learning style objectives can be applied by viewing sampling as a policy rollout and optimizing the induced distribution using policy gradients, often with variants that avoid training an explicit critic (e.g., GRPO-like methods [19]). A common practical choice is to convert deterministic ODE sampling into a stochastic SDE process to inject exploration and to obtain lower-variance updates [8, 9, 11, 23, 26], but this conversion can restrict solver choices and tends to be more effective with longer trajectories. Alternative approaches align diffusion models via direct preference optimization or reward-supervised fine-tuning [22, 25].

Neighbor GRPO [3] provides an alternative perspective by reinterpreting SDE-based GRPO updates as a distance-driven contrastive objective over a neighborhood of candidate trajectories. By constructing neighbors via perturbations of the initial noise and defining a training-time surrogate transition distribution based on softmax distances in trajectory/latent space, it enables informative preference signals while keeping inference-time sampling purely ODE and compatible with high-order solvers. Our method follows this high-level principle—*training-time controllable exploration with inference-time determinism*—but adapts it to the streaming AR video setting. In particular, we formulate alignment at the chunk level to better match autoregressive structure, aiming to provide more localized and stable signals under cross-chunk error propagation, and we introduce a semi-on-policy strategy to improve data reuse and wall-clock efficiency.

3 Methodology

3.1 Preliminary: Neighbor GRPO

Neighbor GRPO [3] aligns flow-based generative models by reinterpreting SDE-based policy optimization as a distance-driven contrastive learning objective, and adopts a contrastive policy optimization (CoPO). To preserve the fast, deterministic inference of ODE sampling, it avoids stochastic SDE conversions. Instead, it constructs a neighborhood of candidate trajectories by perturbing a shared initial noise $\epsilon^* \sim \mathcal{N}(0, I)$:

$$\epsilon^{(i)} = \sqrt{1 - \sigma^2} \epsilon^* + \sigma \delta^{(i)}, \quad \delta^{(i)} \sim \mathcal{N}(0, I), \quad i = 1, \dots, G, \quad (1)$$

where $\sigma \in (0, 1)$ controls the exploration radius. These noises are deterministically rolled out using a reference policy via an ODE solver, and intermediate latents $\{x_t^{(i)}\}_{i=1}^G$ at a chosen timestep t are collected as candidates.

To enable policy gradient updates, Neighbor GRPO defines a surrogate training-time transition distribution based on the distance between an anchor latent $x_t^{(\theta)}$ (produced by the active policy θ from the same timestep) and the candidates:

$$d^{(i)} = \left\| x_t^{(i)} - x_t^{(\theta)} \right\|_2^2, \quad \pi_\theta(i) = \frac{\exp(-d^{(i)}/\tau)}{\sum_{k=1}^G \exp(-d^{(k)}/\tau)}, \quad (2)$$

where τ is a temperature hyperparameter. Given rewards $\{r^{(i)}\}_{i=1}^G$ for the candidates, it computes group-normalized advantages $A^{(i)} = \frac{r^{(i)} - \bar{r}}{\sigma_r}$. Plugging the surrogate policy and advantages into the GRPO objective, the model is optimized to maximize:

$$J(\theta) = \frac{1}{G} \sum_{i=1}^G \min \left(\frac{\pi_\theta(i)}{\pi_{\text{old}}(i)} A^{(i)}, \text{clip} \left(\frac{\pi_\theta(i)}{\pi_{\text{old}}(i)}, 1 - \epsilon, 1 + \epsilon \right) A^{(i)} \right). \quad (3)$$

This objective naturally pulls the anchor towards candidates with positive advantages and pushes it away from those with negative advantages. Furthermore, Neighbor GRPO employs symmetric anchor sampling to reduce training costs by treating multiple candidates as anchors without extra forward passes.

3.2 AR-CoPO

In this section we introduce AR-CoPO. As illustrated in Figure 3, the training pipeline consists of three phases: (1) rolling out neighborhood candidates by forking at a randomly selected pivot chunk, (2) computing sequence-level rewards for each completed branch, and (3) replaying the pivot chunk through the current policy and performing a contrastive GRPO update. We detail each component below.




Fig. 3: The AR-CoPO training pipeline. **(1) Rollout:** The model autoregressively generates a shared context up to a randomly selected pivot chunk p . At chunk p , the base initial noise is perturbed into G neighbors; each neighbor is forked into an independent branch and autoregressively completed to produce a full video sequence. **(2) Reward:** Each completed sequence is decoded and scored by a reward model, yielding a sequence-level reward per branch. **(3) Replay & Update:** The saved pivot-chunk trajectories are replayed through the current policy; distances between current and old $\hat{x}_0$ predictions induce surrogate policy ratios, which are used in a clipped GRPO update confined to the pivot chunk.

Chunk-level Alignment via Forking Due to the streaming nature of autoregressive (AR) generation, naive application of sequence-level GRPO is prohibitively expensive and suffers from severe credit assignment issues. To address this, we propose a chunk-level alignment strategy that performs action space sampling (forking) at a randomly selected chunk, and evaluates the generation via a sequence-level reward.

Specifically, for a target sequence of length L chunks, each optimization iteration proceeds as follows:

- 1. Shared Context Generation:** We randomly sample a pivot chunk index $p \in \{1, \dots, L\}$. The model sequentially generates the first $p - 1$ chunks to establish a shared historical context h_{p-1} (e.g., cached KV states).
- 2. Action Space Forking:** At the p -th chunk, we branch the generation to perform action space sampling. Based on a shared base initial noise ϵ_p^* , we construct a group of G perturbed neighboring noises $\{\epsilon_p^{(i)}\}_{i=1}^G$ following Eq. 1. For each branch i , the model completes the T -step denoising generation to produce the chunk latent $x_p^{(i)}$. The states along this T -step trajectory are stored in a replay buffer.
- 3. Rollout and Sequence-Level Reward:** For each of the G branches, the model deterministically generates the remaining $L - p$ chunks (without fur-

ther perturbation). After the full sequence is completed, we compute a *sequence-level* reward $r^{(i)}$ for each branch.

Controlled noise sharing. Crucially, within each training iteration, the only source of randomness that *differs* across branches is the initial noise of the pivot chunk $\epsilon_p^{(i)}$. All other noise sequences are shared identically across branches: the initial noises of every non-pivot chunk and all CM solver noises at every denoising timestep within every chunk are drawn once and reused across all G branches. This design ensures that the G completed sequences are identical everywhere except in the content generated at chunk p , so any reward difference $r^{(i)} - r^{(j)}$ is cleanly attributable to the choice of $\epsilon_p^{(i)}$ and the resulting latent $x_p^{(i)}$, with no confounding stochasticity from later generation stages.

During the policy update phase, we retrieve the saved trajectories of the p -th chunk from the replay buffer. We compute the advantages $A^{(i)}$ using the sequence-level rewards $r^{(i)}$. Then, utilizing the distance-induced surrogate policy $\pi_\theta(i | s_p)$ (where distances are computed using the chunk latents x_p), we perform a standard Neighbor GRPO parameter update optimizing Eq. 3 restricted to the p -th chunk.

This chunk-level forking mechanism offers two major advantages. First, it significantly reduces training costs: during replay-style gradient updates, backpropagation is strictly confined to the T steps of the single p -th chunk, bypassing the prohibitive cost of full-sequence multi-branch backpropagation. Second, it provides a stable and localized credit assignment signal: by isolating all inter-branch randomness to the initial noise of chunk p while sharing all other noise sources across the group (see above), reward differences among branches are unconfounded and can be directly attributed to the generation choices made at that specific chunk.

CoPO for Consistency Model Alignment The forking mechanism described above provides a general chunk-level alignment framework, but leaves open the question of how to define the surrogate transition distribution $\pi_\theta(i | s_p)$ for different sampler types. Neighbor GRPO was originally designed for flow-matching (FM) models: it measures distances between candidate latents $x_t^{(i)}$ produced by a deterministic ODE solver at an intermediate timestep t (Eq. 2), which is a natural choice when the model is a continuous-time velocity field. For FM-based AR generators, this original distance definition applies directly within our chunk-level framework.

For consistency models (CMs) such as Self-Forcing [5], however, the ODE-solver distance is not well suited. CMs do not follow a standard DDIM/ODE trajectory; their key operation is a one-step mapping from a noisy latent directly to a clean prediction $\hat{x}_0$. Measuring distances in the intermediate x_t space conflates noise scale with semantic content and is therefore less informative. We instead apply the same CoPO principle but define the distance in $\hat{x}_0$ prediction

Algorithm 1 AR-CoPO Training (one iteration)**Require:** Policy θ , reward $r(\cdot)$, sequence length L , group size G

- 1: Sample pivot $p \sim \text{Uniform}(1, L)$
- 2: Generate shared context h_{p-1} by running θ on chunks $1, \dots, p-1$
- 3: **for** $i = 1, \dots, G$ **do** $\triangleright$ Fork at chunk p
- 4: $\epsilon_p^{(i)} \leftarrow \sqrt{1 - \sigma^2} \epsilon_p^* + \sigma \delta^{(i)}, \quad \delta^{(i)} \sim \mathcal{N}(0, I)$
- 5: Denoise chunk p from $\epsilon_p^{(i)}$; complete remaining chunks; compute $r^{(i)}$
- 6: **end for**
- 7: $A^{(i)} \leftarrow (r^{(i)} - \bar{r}) / \sigma_r$
- 8: Replay chunk p : compute $\pi_\theta(i) \propto \exp(-\|\hat{x}_0^{(i)} - \hat{x}_0^{(\theta)}\|^2 / \tau_0)$
- 9: Update θ via GRPO (Eq. 3) on chunk p only




Fig. 4: On-policy *vs.* semi-on-policy training under AR-CoPO. **Left:** On-policy training rolls out fresh candidates from the evolving policy π_θ at each iteration, enabling active exploration of new generation modes guided by the reward signal. **Right:** Semi-on-policy training fixes all rollouts to a reference policy π_{ref} ; the contrastive objective upweights high-reward candidates and suppresses low-reward ones within a trust region maintained by ratio clipping, enhancing exploitation without sacrificing stability. Each paradigm trains an independent LoRA adapter; merging the two adapters yields the final aligned model that benefits from both exploration and exploitation.

space using the CM one-step prediction $\hat{x}_{0,t} = F_\theta(x_t, h_{t-1}, t)$:

$$d_{0,t}^{(i)} = \left\| \hat{x}_{0,t}^{(i)} - \hat{x}_{0,t}^{(\theta)} \right\|_2^2, \quad \pi_\theta(i | s_t) = \frac{\exp(-d_{0,t}^{(i)} / \tau_0)}{\sum_{k=1}^G \exp(-d_{0,t}^{(k)} / \tau_0)}, \quad (4)$$

where $\hat{x}_{0,t}^{(i)} = F_{\theta_{\text{old}}}(x_t^{(i)}, h_{t-1}, t)$ is produced by the old parameters on candidate inputs, $\hat{x}_{0,t}^{(\theta)}$ is the current prediction on the anchor, and τ_0 is the temperature.

3.3 Semi-On-Policy Alignment

Limitations of pure on-policy exploration. The on-policy AR-CoPO framework described above generates diverse candidate rollouts via initial noise perturbation, enabling active exploration of the generation space. However, not all

reward signals respond equally to exploration-driven training. In particular, text alignment (TA)—a global, semantic-level reward that measures whether the generated video faithfully reflects the input prompt—is inherently difficult to improve through local noise perturbation alone. Because TA captures high-level semantic coherence that emerges from the video as a whole, small perturbations to the initial chunk noise typically yield semantically similar outputs with negligible reward variance, providing a weak and noisy gradient signal. This observation is consistent with prior work on aligning bidirectional video generators [9, 26], where training toward TA reward is often omitted altogether due to instability or reward hacking. We argue that the root cause is a fundamental mismatch: exploration-based gradients operate locally in noise space, while semantic quality is a global property of the full sequence that cannot be reliably shaped by random perturbations.

Semi-on-policy training as exploitation. To address this, we complement on-policy exploration with a dedicated *exploitation* paradigm. As illustrated in Figure 4 (right), rather than rolling out from an evolving policy, we fix all rollouts to a reference policy π_{ref} (the initialization checkpoint) and pre-collect a large replay buffer of reference candidates. This reference model already captures a reasonable generation distribution. The contrastive AR-CoPO objective is then applied over these fixed rollouts: high-reward candidates are upweighted while low-reward candidates are suppressed, *without relying on stochastic exploration to discover new generation modes*. Concretely, the replay buffer stores chunk-level entries of the form $(h_{p-1}, \epsilon_p^*, \{\epsilon_p^{(i)}\}_{i=1}^G, \{x_p^{(i)}\}_{i=1}^G, \{r^{(i)}\}_{i=1}^G)$. At each gradient step, we retrieve a batch of these entries, recompute the anchor prediction $\hat{x}_{0,t}^{(\theta)}$ from the current policy, and optimize the AR-CoPO objective in Eq. 3 using the cached candidates and rewards.

Trust region via ratio clipping. A naïve off-policy scheme—applying the contrastive objective to fixed reference rollouts without any constraint—risks distributional shift: the policy may drift far from the reference distribution and collapse. We prevent this by retaining the ratio clipping in the AR-CoPO objective, which enforces a trust region around π_{ref} [17]. This constraint ensures that the policy’s response to any single reference rollout remains bounded, maintaining training stability even when the buffer data becomes stale. We empirically confirm this in Section 4.3: removing ratio clipping causes rapid performance degradation, while the clipped variant preserves generation quality while improving VBench scores.

Combining exploration and exploitation via LoRA merging. The on-policy and semi-on-policy objectives serve complementary purposes and are best optimized independently. We therefore train two separate LoRA adapters—one under on-policy AR-CoPO (for exploration and reward improvement) and one under semi-on-policy AR-CoPO (for exploitation and overall quality)—and merge them at inference time. As shown in Figure 4, this separation avoids interference between the two objectives during training. The merged model benefits from both paradigms: the semi-on-policy adapter reshapes the in-distribution qual-

Table 1: Quantitative comparison of streaming AR video generation models on VBench (Quality, Semantic, Total) and VideoAlign (VQ, MQ, TA, Overall) benchmarks. “+ ours (Semi)” denotes the model after semi-on-policy alignment; “+ ours (Merged)” denotes the final model obtained by merging the semi-on-policy and on-policy LoRA adapters.

Method	VBench			VideoAlign			
	Quality	Semantic	Total	VQ	MQ	TA	Overall
Self-Forcing	84.87	71.27	82.15	3.80	1.68	2.28	7.76
Causal-Forcing	85.27	70.35	82.28	3.97	1.43	2.40	7.79
LongLive	85.10	71.16	82.31	3.87	1.76	2.43	8.06
Self-Forcing	84.87	71.27	82.15	3.80	1.68	2.28	7.76
+ ours (semi)	85.15	71.68	82.45	3.70	1.60	2.30	7.61
+ ours (on-policy)	84.81	70.71	81.99	4.15	2.06	2.30	8.51
+ ours (merged)	85.07	70.55	82.17	4.00	1.86	2.36	8.22

ity of the reference policy, while the on-policy adapter steers the model toward higher reward regions through active exploration.

4 Experiments

We evaluate AR-CoPO on Self-Forcing [5], which is the baseline model of many variants. We also evaluate Causal-Forcing [33] and report the results in the supplementary material. We further report representative few-step streaming AR video generators as strong baselines such as LongLive [27]. Training is conducted on MovieGen Video Bench [15]. We optimize toward the VideoAlign [12] reward suite, which comprises text alignment (TA), video quality (VQ), and motion quality (MQ). For the main experiments, all three rewards (TA, MQ, VQ) are jointly optimized. For holistic evaluation, we additionally report VBench scores (Quality, Semantic, Total). All models are fine-tuned with LoRA (rank 64, $\alpha = 128$) using 24 GPUs. We use a group size of $G = 12$ and a learning rate of 1×10^{-5} . Following Neighbor GRPO [3], we adopt symmetric anchor sampling with an anchor batch size of 4. For the semi-on-policy strategy, we collect a replay buffer of 100 rollout groups from the initialization model.

4.1 Comparison with AR Models

Quantitative results. Table 1 reports the main quantitative comparison. Semi-on-policy training alone surpasses all streaming AR baselines on VBench Total (82.45 vs. 82.31 for LongLive), with consistent gains on both Quality and Semantic dimensions, demonstrating the effectiveness of exploitation-focused alignment. After merging the on-policy LoRA adapter, VideoAlign Overall improves from 7.76 to 8.22. Crucially, this improvement is accompanied by a maintained VBench Total (82.15→82.17), confirming that the gain reflects genuine alignment rather than in-domain score inflation. We elaborate on why this dual-benchmark criterion is the appropriate measure of alignment quality in Section 4.4.




Fig. 5: Qualitative comparison between AR-CoPO (up) and Self-Forcing (down) on diverse text prompts. AR-CoPO produces videos with improved visual fidelity, motion quality, and better adherence to the text prompt.

Qualitative results. Figure 5 shows side-by-side frame comparisons between AR-CoPO and Self-Forcing across diverse prompts. AR-CoPO produces videos with better aesthetic quality, more vivid appearance, more coherent motion, and better adherence to the text description.

4.2 Comparison with SDE-GRPO

We compare AR-CoPO against an SDE-based GRPO baseline following the design of Dance-GRPO [26] and FlowGRPO [11]. Training curves are shown in Figure 2 and Figure 7. The SDE-based variant fails to improve the reward throughout training, whereas AR-CoPO steadily achieves higher scores.

We attribute this failure to the near-deterministic nature of distilled few-step consistency models (CMs). Although CM solvers inject intermediate re-noise during inference, sample diversity in Self-Forcing [5] is almost entirely governed by the *initial* chunk noise rather than the intermediate solver noise. In other words, the model behaves as an approximately deterministic mapping from initial noise to clean frames, with intermediate stochasticity playing a negligible role.

To validate this, we conduct a controlled noise substitution study (Figure 6): we freeze all random draws in the sampler and selectively replace individual




Fig. 7: VQ training curves for AR-CoPO (ours) and the SDE-based GRPO baseline.




Fig. 6: Analysis of entropy sources in Self-Forcing. Each sub-figure corresponds to forking at a different chunk position. **Row 1:** Reference sample with all noise frozen. **Row 2:** Only the initial noise of the forked chunk is replaced—the output changes substantially. **Rows 3–5:** Only the CM solver noise at a specific denoising timestep within the chunk is replaced—the output changes marginally. This confirms that sample diversity in Self-Forcing is governed almost entirely by the initial noise, making intermediate SDE-style noise injection ineffective as an exploration mechanism.

noise tensors. Replacing only the initial chunk noise (Row 2) causes substantial variation in the output, while replacing the intermediate CM solver noise at any individual denoising step (Rows 3–5) produces almost no visible change. Crucially, this effect is more pronounced in later chunks, where the accumulated autoregressive context further constrains the generation and makes it even more sensitive to the initial noise.

SDE-based GRPO methods [8, 9, 11, 26] such as DanceGRPO typically freeze the initial noise and define the action space over intermediate SDE noise injections. Because these intermediate injections carry negligible entropy in few-step AR models, the resulting policy gradient signal is near-zero, explaining why SDE-based training cannot drive reward improvement in this setting.

4.3 Effect of Semi-On-Policy Training

To isolate the effect of each training paradigm on semantic alignment, we ablate three training strategies under the AR-CoPO objective optimizing *only* the TA reward: on-policy, semi-on-policy, and fully off-policy (without ratio clipping).

Table 2: Ablation of training strategies on VideoAlign and VBench metrics.

Method	VBench			VideoAlign			
	Quality	Semantic	Total	VQ	MQ	TA	Overall
Self-Forcing	84.87	71.27	82.15	3.80	1.68	2.28	7.76
on-policy	81.66	69.68	79.26	3.53	0.25	2.63	6.42
off-policy	69.78	60.84	67.99	2.22	-0.15	2.16	4.23
semi-on-policy	85.15	71.68	82.45	3.70	1.60	2.30	7.61

**Fig. 8:** Comparison of generated samples under semi-on-policy (top for each) and on-policy (bottom for each) training. On-policy training introduces visible temporal inconsistencies, whereas semi-on-policy training preserves generation quality.

This controlled setting allows us to directly assess whether semi-on-policy training can improve TA without the confounding effects of simultaneously optimizing VQ and MQ. Results are summarized in Table 2.

On-policy training improves the in-domain TA score, but causes severe degradation across all other metrics and introduces visible generation artifacts (see Figure 8). Most strikingly, MQ collapses from 1.68 to 0.25, accompanied by a sharp drop in VBench Total (82.15→79.26). We identify this as a core failure mode of on-policy TA optimization. TA is a global, semantic-level reward: it measures whether the generated video as a whole faithfully reflects the input prompt. Optimizing such a signal via local noise-space exploration drives the model toward semantic shortcuts, *i.e.* reward hacking, at the expense of temporal coherence and structural continuity. This manifests as the severe motion discontinuities and frame inconsistencies shown in Figure 8, which directly cause MQ and VQ to collapse. This instability is consistent with prior work [9, 26], where TA optimization is often omitted altogether due to similar failure modes.

Semi-on-policy training avoids this collapse: VideoAlign scores remain broadly on par with the Self-Forcing baseline, while VBench Quality and Semantic scores improve over it. The ratio-clipping constraint is essential here—it keeps the policy within a trust region of the reference rollouts and prevents out-of-distribution drift. We confirm this by comparing against **fully off-policy**

Table 3: Scores when merging on-policy weights with different strength

Scale	VBench			VideoAlign			
	Quality	Semantic	Total	VQ	MQ	TA	Overall
1.0	84.90	70.38	81.99	4.13	1.86	2.34	8.33
0.8	85.07	70.55	82.17	4.00	1.86	2.36	8.22
0.6	85.11	70.72	82.23	3.86	1.78	2.36	7.99
0.4	85.14	71.44	82.40	3.76	1.62	2.34	7.72
0 (Semi)	85.15	71.68	82.45	3.70	1.60	2.30	7.61

training (no ratio clipping): without the trust-region constraint, the model drifts far from the reference distribution, causing most scores to deteriorate.

4.4 LoRA Merging

The on-policy and semi-on-policy adapters are trained independently and merged at inference time by scaling the on-policy LoRA weights before addition. As shown in Table 3, varying the scale reveals a clear monotonic trade-off: increasing the on-policy contribution monotonically improves VideoAlign Overall (7.61→8.33) while simultaneously degrading VBench Total (82.45→81.99), and vice versa. This reflects the complementary nature of the two adapters—the semi-on-policy adapter exploits high-quality reference rollouts to preserve generation quality, whereas the on-policy adapter drives reward-seeking exploration that improves human preference scores.

Scale selection criterion. As discussed in Section 4.3, directly optimizing AR video generators toward reward models is prone to in-domain reward hacking; the on-policy TA experiment demonstrates that a higher reward score does not always indicate better generation quality. Consequently, a higher VideoAlign Overall score alone is an unreliable criterion for model selection. We instead adopt a stricter criterion: the selected configuration must improve *both* the in-domain benchmark (VideoAlign Overall) and the independent out-of-domain benchmark (VBench Total) relative to the baseline. Among all scales in Table 3, scale= 1.0 achieves the highest VideoAlign Overall (8.33) but at the cost of VBench Total degradation (81.99 vs. 82.15 for the baseline), which we interpret as a sign of over-optimization rather than genuine quality improvement. Scale= 0.8 is the largest scale that satisfies the dual-improvement criterion—VideoAlign Overall rises from 7.76 to 8.22 while VBench Total is maintained (82.15→82.17)—and is therefore selected as the default reported setting.

5 Conclusion

In this paper, we present AR-CoPO, a chunk-level contrastive policy optimization framework for aligning few-step streaming autoregressive video generators to human preference. By constructing neighborhood candidates through a forking

mechanism at a randomly selected pivot chunk, AR-CoPO circumvents the fundamental mismatch between SDE-based exploration and the near-deterministic dynamics of consistency model samplers. The complementary semi-on-policy strategy further improves generation quality by exploiting high-quality reference rollouts within a trust region, without sacrificing the exploration benefits of on-policy training. Experiments on Self-Forcing demonstrate consistent improvements on both VBench and VideoAlign, validating the effectiveness of AR-CoPO for post-training alignment of streaming AR video generation.

Acknowledgements

This work is supported in part by NSFC-RGC Project N_CUHK498/24, in part by Guangdong Basic and Applied Basic Research Foundation (No. 2023B1515130008, XW), in part by Shenzhen Loop Area Institute, and in part by the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission (ITC)’s InnoHK.

References

1. Christiano, P.F., Leike, J., Brown, T., Martic, M., Legg, S., Amodei, D.: Deep reinforcement learning from human preferences. *Advances in neural information processing systems* **30** (2017)
2. Ge, X., Zhang, Y., Huang, Y., He, D., Wang, X., Ma, B., Song, G., Liu, Y., Zhang, J.: Salt: Self-consistent distribution matching with cache-aware training for fast video generation. *arXiv preprint arXiv:2604.03118* (2026)
3. He, D., Feng, G., Ge, X., Niu, Y., Zhang, Y., Ma, B., Song, G., Liu, Y., Li, H.: Neighbor grpo: Contrastive ode policy optimization aligns flow models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6033–6042 (2026)
4. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
5. Huang, X., Li, Z., He, G., Zhou, M., Shechtman, E.: Self forcing: Bridging the train-test gap in autoregressive video diffusion. *arXiv preprint arXiv:2506.08009* (2025)
6. 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)
7. Labs, B.F., Batifol, S., Blattmann, A., Boesel, F., Consul, S., Diagne, C., Dockhorn, T., English, J., English, Z., Esser, P., Kulal, S., Lacey, K., Levi, Y., Li, C., Lorenz, D., Müller, J., Podell, D., Rombach, R., Saini, H., Sauer, A., Smith, L.: Flux.1 kontext: Flow matching for in-context image generation and editing in latent space (2025), <https://arxiv.org/abs/2506.15742>
8. Li, J., Cui, Y., Huang, T., Ma, Y., Fan, C., Yang, M., Zhong, Z.: Mixgrpo: Unlocking flow-based grpo efficiency with mixed ode-sde. *arXiv preprint arXiv:2507.21802* (2025)
9. Li, Y., Wang, Y., Zhu, Y., Zhao, Z., Lu, M., She, Q., Zhang, S.: Branchgrpo: Stable and efficient grpo with structured branching in diffusion models. *arXiv preprint arXiv:2509.06040* (2025)

10. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. arXiv preprint arXiv:2210.02747 (2022)
11. Liu, J., Liu, G., Liang, J., Li, Y., Liu, J., Wang, X., Wan, P., Zhang, D., Ouyang, W.: Flow-grpo: Training flow matching models via online rl. arXiv preprint arXiv:2505.05470 (2025)
12. Liu, J., Liu, G., Liang, J., Yuan, Z., Liu, X., Zheng, M., Wu, X., Wang, Q., Qin, W., Xia, M., et al.: Improving video generation with human feedback. arXiv preprint arXiv:2501.13918 (2025)
13. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in neural information processing systems* **35**, 5775–5787 (2022)
14. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. *Advances in neural information processing systems* **35**, 27730–27744 (2022)
15. Polyak, A., Zohar, A., Brown, A., Tjandra, A., Sinha, A., Lee, A., Vyas, A., Shi, B., Ma, C.Y., Chuang, C.Y., et al.: Movie gen: A cast of media foundation models. arXiv preprint arXiv:2410.13720 (2024)
16. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
17. Schulman, J., Levine, S., Abbeel, P., Jordan, M., Moritz, P.: Trust region policy optimization. In: *International conference on machine learning*. pp. 1889–1897. PMLR (2015)
18. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017)
19. Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang, H., Zhang, M., Li, Y., Wu, Y., et al.: Deepseekmath: Pushing the limits of mathematical reasoning in open language models. arXiv preprint arXiv:2402.03300 (2024)
20. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
21. Stiennon, N., Ouyang, L., Wu, J., Ziegler, D., Lowe, R., Voss, C., Radford, A., Amodei, D., Christiano, P.F.: Learning to summarize with human feedback. *Advances in neural information processing systems* **33**, 3008–3021 (2020)
22. Wallace, B., Dang, M., Rafailov, R., Zhou, L., Lou, A., Purushwalkam, S., Ermon, S., Xiong, C., Joty, S., Naik, N.: Diffusion model alignment using direct preference optimization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8228–8238 (2024)
23. Wang, Y., Li, Z., Zang, Y., Zhou, Y., Bu, J., Wang, C., Lu, Q., Jin, C., Wang, J.: Pref-grpo: Pairwise preference reward-based grpo for stable text-to-image reinforcement learning. arXiv preprint arXiv:2508.20751 (2025)
24. Wu, B., Zou, C., Li, C., Huang, D., Yang, F., Tan, H., Peng, J., Wu, J., Xiong, J., Jiang, J., et al.: Hunyuanvideo 1.5 technical report. arXiv preprint arXiv:2511.18870 (2025)
25. Wu, X., Hao, Y., Zhang, M., Sun, K., Huang, Z., Song, G., Liu, Y., Li, H.: Deep reward supervisions for tuning text-to-image diffusion models. In: *European Conference on Computer Vision*. pp. 108–124. Springer (2024)
26. Xue, Z., Wu, J., Gao, Y., Kong, F., Zhu, L., Chen, M., Liu, Z., Liu, W., Guo, Q., Huang, W., et al.: Dancegrpo: Unleashing grpo on visual generation. arXiv preprint arXiv:2505.07818 (2025)

27. Yang, S., Huang, W., Chu, R., Xiao, Y., Zhao, Y., Wang, X., Li, M., Xie, E., Chen, Y., Lu, Y., Chen, S.H.Y.: Longlive: Real-time interactive long video generation (2025)
28. Yin, T., Gharbi, M., Park, T., Zhang, R., Shechtman, E., Durand, F., Freeman, W.T.: Improved distribution matching distillation for fast image synthesis. In: NeurIPS (2024)
29. Yin, T., Gharbi, M., Zhang, R., Shechtman, E., Durand, F., Freeman, W.T., Park, T.: One-step diffusion with distribution matching distillation. In: CVPR (2024)
30. 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: CVPR (2025)
31. Zhang, Y., Yang, H., Zhang, Y., Hu, Y., Zhu, F., Lin, C., Mei, X., Jiang, Y., Peng, B., Yuan, Z.: Waver: Wave your way to lifelike video generation. arXiv preprint arXiv:2508.15761 (2025)
32. Zheng, Z., Peng, X., Yang, T., Shen, C., Li, S., Liu, H., Zhou, Y., Li, T., You, Y.: Open-sora: Democratizing efficient video production for all. arXiv preprint arXiv:2412.20404 (2024)
33. 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. arXiv preprint arXiv:2602.02214 (2026)