

Salt: Self-Consistent Distribution Matching with Cache-Aware Training for Fast Video Generation

Xingtong Ge^{1,2}, Yi Zhang², Yushi Huang¹, Dailan He², Xiahong Wang²,
Bingqi Ma², Guanglu Song², Yu Liu², and Jun Zhang¹^{*}

¹ Hong Kong University of Science and Technology

² Vivix Group Limited

xingtong.ge@gmail.com, eejzhang@ust.hk

Abstract. Distilling video generation models to extremely low inference budgets (e.g., 2–4 NFEs) is crucial for real-time deployment, yet remains challenging. Trajectory-style consistency distillation often becomes conservative under complex video dynamics, yielding an over-smoothed appearance and weak motion. Distribution matching distillation (DMD) can recover sharp, mode-seeking samples, but its local training signals do not explicitly regularize how denoising updates compose across timesteps, making composed rollouts prone to drift. To overcome this challenge, we propose Self-Consistent Distribution Matching Distillation (SC-DMD), which explicitly regularizes the endpoint-consistent composition of consecutive denoising updates. For real-time autoregressive video generation, we further treat the KV cache as a quality parameterized condition and propose Cache-Distribution-Aware training. This training scheme applies SC-DMD over multi-step rollouts and introduces a cache-conditioned feature alignment objective that steers low-quality outputs toward high-quality references. Across extensive experiments on both non-autoregressive backbones (e.g., Wan 2.1) and autoregressive real-time paradigms (e.g., Self Forcing), our method, dubbed **Salt**, consistently improves low-NFE video generation quality while remaining compatible with diverse KV-cache memory mechanisms. Project page: <https://xingtongge.github.io/Salt>.

Keywords: Video generation · Diffusion distillation · Autoregressive generation

1 Introduction

Despite rapid progress in diffusion and flow-matching models [12, 22, 25, 35], pushing video generation to extremely low inference budgets (e.g., 2–4 NFEs) remains challenging. A natural strategy, inherited from few-step image generation, is to distill a multi-step teacher into a few-step student using trajectory-level consistency objectives [3, 7, 9, 28, 33]. However, video dynamics are inherently more complex, as the same conditioning may admit multiple plausible futures—for

^{*} Corresponding Author

example, the continuation of a motion across time can vary in speed, extent, or trajectory. Under regression-style trajectory losses, this ambiguity tends to pull the predictor toward conditional averages, often resulting in over-smoothed appearance and conservative motion [10, 30, 44].

By contrast, distribution-level score distillation—notably Distribution Matching Distillation (DMD) [38, 41, 42]—has emerged as a widely adopted paradigm for practical video distillation. By matching the student’s output distribution to a strong reference through adversarial or distribution-matching gradients, DMD avoids regressing toward a single target trajectory and often produces sharper, more expressive samples in the low-NFE regime. This advantage has driven rapid adoption across recent video distillation methods [5, 6, 14, 15, 20, 21, 31]. It has also become a core ingredient in autoregressive real-time generation systems—including CausVid [43], Self Forcing [13], and LongLive [39]—suggesting that DMD-style objectives have become a practical default for few-step video generation. Crucially, most of these systems operate at 2–8 denoising steps rather than a single step, meaning that inference quality depends on repeatedly composing the learned denoising operator across multiple steps.

However, this reliance on multi-step composition exposes a structural gap in the DMD framework. Few-step sampling is fundamentally a problem of repeatedly composing learned denoising updates, which is actually incompatible with the objective of predicting clean outputs in DMD training. As a result, per-step errors can accumulate along the composed trajectory, leading to drift as the number of steps increases [32]. In this work, we note that the training objective for DMD is purely local: each noise level is supervised independently as a “one-step generator”, with no constraint on how the per-timestep behaviors interact when applied in sequence. We term this *compositionality deficit*: despite strong single-step quality, increasing the number of denoising steps can degrade rather than improve output (e.g., over-exposure and loss of fine-grained semantics; see Fig. 1). This issue may be further amplified in autoregressive settings [13, 39, 48], where each newly generated chunk is conditioned on previous generations; errors introduced in earlier chunks can be encoded into the generated KV cache and progressively propagated to later chunks over time.

In this work, we propose **Self-Consistent Distribution Matching Distillation (SC-DMD)**. The key idea is to retain DMD as the primary driver of distribution alignment and single-step sharpness, while augmenting it with a *shortcut self-consistency regularizer* to mitigate the compositionality deficit. Concretely, this regularizer penalizes the semigroup defect of the student-induced transport operator, encouraging endpoint consistency under the composition of consecutive student updates. Rather than replacing distribution matching with a trajectory-style objective, we use this constraint as a complementary structural bias: the self-consistency term improves compositionality, while DMD preserves sharpness and avoids the mode-averaging tendencies of pure trajectory-based supervision in video. This combination yields consistent improvements in both non-autoregressive video generation (e.g., Wan 2.1 [35]) and autoregressive real-time generation. Moreover, autoregressive real-time video generation

introduces an additional challenge beyond few-step sampling. With KV caching, each newly generated chunk is conditioned on previous generations, and the quality of this conditioning depends on how accurately earlier chunks were denoised; consequently, errors can propagate through the cache and accumulate over long rollouts [13, 43]. We find that training with a single fixed step count leaves the generator brittle to the cache-quality variation encountered at inference. To address this, we introduce **Cache-Aware Mixed-Step Training** for autoregressive generation. Specifically, training alternates among rollouts with $K \in \{2, 4, 8\}$ denoising steps, exposing the model to cache-conditioned inputs of varying quality and significantly improving few-step generation with the application of self-consistency mechanism. Besides, this mixed-step setup naturally enables cross-step supervision: under the same noise and context, higher-step outputs can be used to regularize lower-step outputs through *cache-conditioned feature alignment*. By transferring richer information from high-step predictions to low-step ones, this design further improves extreme low-step performance and rollout stability. In summary, our contributions are threefold:

- We identify the *compositionality deficit* of DMD: although distribution-level score distillation is widely used, its purely local training signal imposes no structural constraint on multi-step coherence.
- We propose SC-DMD, which addresses the compositionality deficit by introducing a semigroup-defect regularizer. It improves multi-step stability while preserving the sharpness of distribution matching, yielding consistent gains in both non-autoregressive and autoregressive video generation.
- We further propose a cache-aware mixed-step training strategy for autoregressive distillation, combining mixed-step rollout with self-consistency and cache-conditioned representation alignment. This further improves extreme low-step generation and long-horizon rollout stability.

Across extensive experiments on both non-autoregressive (e.g., Wan 2.1) and autoregressive real-time paradigms (e.g., Self Forcing and LongLive), our full method, dubbed **Salt**, consistently improves extremely low-NFE video generation quality and long-horizon stability in autoregressive rollouts, while retaining compatibility with diverse KV-cache memory mechanisms [13, 27, 39] and incurring no additional inference overhead.

2 Related Work

Diffusion Distillation methods mainly fall into two categories: trajectory-based [26, 28, 29, 32, 33, 36] and distribution-based approaches [5, 6, 8, 11, 16, 20, 31, 37, 41, 42]. Trajectory-based distillation compresses sampling by approximating the teacher’s denoising dynamics: DCM [29] mitigates the detail degradation by decoupling semantic learning from detail refinement via a dual-expert design. rCM [47] scales continuous-time consistency distillation to application-scale image and video diffusion through efficient JVP training, and incorporates score-distillation as a long-skip regularizer to preserve fine details under few steps

generation. Distribution-based methods instead target matching the generative distribution: Recent one-step video methods [20, 21] operationalize distribution matching via adversarial post-training, while POSE [5] introduces a phased equilibrium procedure to stabilize adversarial one-step distillation for video models. **Autoregressive Video Generation** enables streaming and interactivity via causal factorization, and can be grouped into training-based [4, 13, 24, 27, 34, 39, 43] and training-free [40, 46] approaches. Training-based methods close the train-test gap of long rollouts beyond standard teacher forcing: Self Forcing [13] directly trains on autoregressive self-rollouts so each step conditions on previously generated context, reducing exposure bias; Causal Forcing [48] further pinpoints the bidirectional-to-causal distillation gap and uses an autoregressive teacher for ODE initialization before asymmetric DMD [42], leading to higher-quality real-time interactive generation; Reward Forcing [27] distills bidirectional video diffusion into a few step autoregressive student with rewarded distribution matching to improve efficient streaming generation. In addition, long-horizon AR generation emphasize scalable causal designs and long-video interaction mechanisms: MAGI-1 [34] scales chunk-wise autoregressive denoising for strong temporal consistency and deployability, LongLive [39] targets real-time interactive long videos with mechanisms for stable prompt transitions and long-range consistency under causal decoding. On the other hand, training-free methods, such as Infinity-RoPE [40] modifies temporal encoding and KV-cache behaviors to unlock action-controllable infinite self-rollout, while UltraViCo [46] studies length extrapolation failure modes and suppresses out-of-window attention to reduce repetition and quality degradation.

3 Method

3.1 Preliminaries

Diffusion/Flow-Matching Models and Flow Maps. Let $x_t \in \mathbb{R}^d$ denote a noisy latent at noise level $t \in [0, 1]$ (larger t indicates noisier states), and c denote conditioning (e.g., text, image, or autoregressive context). Many diffusion and flow-matching models admit a *probability flow ODE* view, where sampling corresponds to transporting states from high noise to low noise via a learned vector field:

$$\frac{dx_t}{dt} = v(x_t, t, c), \quad t \in [0, 1], \quad x_{t_0} \sim p(x_{t_0}). \quad (1)$$

The induced timestep-to-timestep transport operator (flow map) deterministically maps a state at timestep t to timestep $s < t$:

$$x_s = \Phi^{t \rightarrow s}(x_t; c), \quad 0 \leq s < t \leq 1. \quad (2)$$

A fundamental property of exact ODE flow maps is the *composition (semigroup) law*:

$$\Phi^{t \rightarrow s} = \Phi^{u \rightarrow s} \circ \Phi^{t \rightarrow u}, \quad \forall t > u > s, \quad (3)$$

which has been explicitly exploited by recent flow-map learning and distillation frameworks [1, 2, 19, 32]. In practice, Φ is approximated by a numerical solver (e.g., Euler) using a learned vector field v_θ .

Distribution Matching Distillation (DMD). Assume a pre-trained teacher diffusion (or flow-matching) model is available, inducing a reference distribution $p_r(x_t)$ and a teacher score $s_r(x_t, t) = \nabla_{x_t} \log p_r(x_t)$. Distribution Matching Distillation (DMD) distills a few-step generator by matching the student distribution to the teacher distribution at sampled noise levels. Instead of directly optimizing $D_{\text{KL}}(p_g \| p_r)$, DMD introduces an intermediate “fake” distribution $p_f(x_t, t)$ parameterized by a learnable score $s_\phi(x_t, t) = \nabla_{x_t} \log p_f(x_t, t)$, and alternates between updating the critic (ϕ) and the generator (θ).

In practice, the generator produces a clean sample $x_0 = G_\theta(z)$ with $z \sim \mathcal{N}(0, I)$, and a noisy sample at level t is obtained by forward diffusion $x_t \sim q(x_t | x_0)$ (e.g., $x_t = \alpha(t)x_0 + \sigma(t)\epsilon$, $\epsilon \sim \mathcal{N}(0, I)$), where x_t depends on θ through x_0 . DMD yields a generator update direction of the form

$$\nabla_\theta \mathcal{L}_{\text{DMD}} = \mathbb{E}_{t,z,\epsilon} \left[(s_\phi(x_t, t) - s_r(x_t, t)) \frac{\partial x_t}{\partial \theta} \right], \quad (4)$$

3.2 SC-DMD: Self-Consistent Distribution Matching Distillation

We distill a pre-trained teacher into a few-step student generator parameterized by a velocity field $v_\theta(x_t, t, c)$ under the probability flow ODE view. The student induces a one-step Euler transport operator

$$\Psi_\theta^{t \rightarrow s}(x) \triangleq x - (t - s) v_\theta(x, t, c), \quad 0 \leq s < t \leq 1, \quad (5)$$

and a K -step Euler sampler on a schedule $\mathcal{T}^{(K)} = \{t_0 > \dots > t_K\}$ iterates $x_{t_{i+1}} = \Psi_\theta^{t_i \rightarrow t_{i+1}}(x_{t_i})$. At inference, we consider both deterministic Euler integration and consistency-style sampling (CM solver) [33], which alternates denoising and re-noising and is widely used in modern autoregressive generation settings [13, 43]. Since both solvers repeatedly apply the same student model, generation quality ultimately depends not only on the quality of each individual update, but also on how well these updates compose across timesteps.

Compositionality deficit of DMD. DMD provides *local* distribution-matching supervision at sampled noise levels (Eq. 4), encouraging each trained timestep to act as a strong one-step generator. However, few-step inference is inherently compositional: it relies on chaining the learned update operators across timesteps. Even in recent DMD variants [13, 41] that use backward simulation to better approximate inference-time inputs, the training objective itself remains *per-timestep and independent across noise levels*. At timestep t_i , DMD only supervises the output distribution of the current step given input x_{t_i} ; it does not explicitly penalize how neighboring updates interact under composition. As a result, the learned operators can be individually well-calibrated yet collectively incoherent, leading to step-count-dependent behavior and unstable multi-step rollouts.



Fig. 1: Compositionality deficit of DMD. First, middle, and last frames from 4-/8-/16-step DMD students (rows, top to bottom) on: (a) “...a spaceman wearing a red wool knitted motorcycle helmet...” and (b) “...a large stack of vintage televisions all showing different programs...museum gallery.” Increasing the number of denoising steps degrades rather than improves quality: the 16-step model loses the knitted helmet texture and corrupts the motorcycle structure, and produces incoherent television details.

We empirically verify this failure mode in a clean, non-autoregressive setting. Using the Wan 2.1 1.3B T2V backbone [35], we train three DMD students on 4-, 8-, and 16-timestep grids, respectively, each receiving distribution-matching supervision on its own set of noise levels using the standard backward-simulation strategy. We then compare generation quality across these step counts using the same inputs [33]. As shown in Fig. 1, increasing the number of denoising steps does not improve generation quality; instead, longer rollouts exhibit artifacts such as over-exposure and semantic degradation. This mirrors the degradation that AYF [32] formally proves can arise in consistency models under the same CM solver. Our observation suggests that the issue is more fundamental: it stems from the absence of explicit compositional constraints in the learned denoising operators, rather than from a particular distillation objective.

SC-DMD: shortcut self-consistency as semigroup-defect regularization. To address this issue, we augment DMD with a shortcut self-consistency regularizer. For an exact flow map $\Phi^{t \rightarrow s}$, the semigroup law in Eq. 3 holds. We therefore encourage the student-induced Euler operator Ψ_θ to approximately satisfy the same compositional property on the training grid. Given a triple (t_s, t_m, t_e) with $t_s > t_m > t_e$, we compute a *direct (one-step) endpoint*

$$x_{t_e}^{(1)} = \Psi_\theta^{t_s \rightarrow t_e}(x_{t_s}) = x_{t_s} - (t_s - t_e) v_\theta(x_{t_s}, t_s, c), \quad (6)$$

and a *composed (two-step) endpoint*

$$x_{t_e}^{(2)} = \Psi_\theta^{t_m \rightarrow t_e}(\Psi_\theta^{t_s \rightarrow t_m}(x_{t_s})). \quad (7)$$

The self-consistency (SC) loss then penalizes the discrepancy between these two endpoints:

$$\mathcal{L}_{\text{SC}} = \mathbb{E} \left[d(x_{t_e}^{(1)}, x_{t_e}^{(2)}) \right], \quad d(a, b) = \|a - b\|_2^2. \quad (8)$$

Intuitively, this regularizer enforces that a direct update from t_s to t_e agrees with the composed update $t_s \rightarrow t_m \rightarrow t_e$, thereby explicitly correcting the compositionality deficit identified above.




Fig. 2: Comparison of training trajectories for few-step distillation methods. Consistency distillation, shortcut models, and flow-map distillation all impose composition-related constraints in different forms. SC-DMD addresses the compositionality deficit by introducing a semigroup-defect regularizer that aligns direct and composed updates while preserving DMD as the distribution matching objective.

Grid and sampling design. Vanilla DMD typically trains on a coarse grid aligned with the K -step inference schedule, leaving little room to sample intermediate times t_m . We instead train on a finer grid $\mathcal{T}_{\text{train}}$ (e.g., $|\mathcal{T}_{\text{train}}|=8$) to expose the model to richer intermediate-step structure. For each $t_s \in \mathcal{T}_{\text{train}}$, we sample the shortcut triple as

$$t_e \sim \mathcal{T}^{(K)} \cap \{t : t < t_s\}, \quad t_m \sim \mathcal{T}_{\text{train}} \cap (t_e, t_s), \quad (9)$$

anchoring t_e to the inference grid so that the shortcut constraint directly supervises the endpoints that matter at test time, while drawing t_m from the finer training grid where dense DMD supervision is available.

Finally, we train the student with

$$\min_{\theta} \mathcal{L}_{\text{DMD}}(\theta; \psi) + \lambda_{\text{SC}} \mathcal{L}_{\text{SC}}(\theta), \quad (10)$$

where the critic parameters ψ are updated by the standard DMD alternating procedure. The SC term adds only lightweight overhead—one additional forward pass at t_m —yet explicitly couples the student’s per-step operators into a more coherent multi-step trajectory. As a result, it directly targets the failure mode in Fig. 1. Ablations in Sec. 4.4 confirm this effect: while vanilla DMD-8 suffers from quality degradation and dynamic collapse, SC-DMD achieves robust improvements on both metrics from 4- to 8-step inference.

Connection to shortcut models and flow map distillation. The self-consistency loss in Eq. (8) shares a common mathematical root with shortcut models [7] and continuous-time flow map distillation [32]: all of them encourage approximate semigroup composition of learned transport operators. As illustrated in Fig. 2, shortcut models parameterize the step size explicitly and use composition consistency as a self-distillation target, while AYT [32] enforces the semigroup law across continuous (t, s) pairs as the primary distillation objective. In contrast, SC-DMD uses the semigroup-defect penalty as a *regularizer* inside a distribution-matching framework. This distinction is important: when semigroup-style constraints are used as the sole training signal, they reduce to regression losses that can induce mode averaging under video dynamics. By retaining DMD as the primary driver of distribution alignment, SC-DMD preserves mode-seeking fidelity




Fig. 3: Overview of Salt for autoregressive video generation. *Left:* A step count $K \in \{8, 4, 2\}$ is sampled to define the few-step denoising trajectory. *Middle:* Conditioned on the current KV cache, text, and noise, the student generator G_θ denoises the current chunk. A self-consistency (SC) loss $\mathcal{L}_{SC}$ regularizes the endpoint discrepancy between a direct update and a composed two-step update. *Right:* During mixed-step autoregressive training, the predicted clean sample is re-noised at a random timestep and optimized with the DMD loss $\mathcal{L}_{DMD}$, while a higher-quality reference branch provides the cache-conditioned alignment loss $\mathcal{L}_{align}$.

while using compositional regularization to rectify the compositionality deficit. We further analyze this design choice in the supplementary material, showing that the SC and DMD losses should act on an aligned training grid to avoid conflicting supervision.

3.3 Cache-condition-aware Training for AR Distillation

In autoregressive video generation, each newly generated chunk is conditioned on the KV cache of previously generated chunks. Unlike standard conditioning inputs (*e.g.* text), the KV cache is itself model-generated, and its quality varies with the quality of prior rollouts. In practice, higher-step rollouts tend to produce richer and more faithful caches, whereas lower-step rollouts produce noisier and more degraded ones. As a result, autoregressive inference exposes the model to a range of cache qualities rather than a single conditioning regime.

This view reveals two practical challenges. First, if training always uses a fixed step count K , the model only observes one cache-quality regime, leading to a train–inference mismatch when rollout quality changes. Second, as autoregressive errors accumulate across chunks, the KV cache can progressively drift away from the conditioning states seen during training, causing semantic and quality degradation.

We address both challenges under a single principle: *systematically cover the range of cache qualities encountered during autoregressive training, while selectively imposing stronger regularization where long-horizon composition is most fragile*. Concretely, we introduce two training mechanisms: (i) mixed-step rollout, which exposes the model to diverse cache qualities and enables SC-DMD in the autoregressive setting, and (ii) cache-conditioned reference alignment, which regularizes outputs across cache-quality levels.

Mixed-step rollout: covering cache quality and enabling SC-DMD. To expose the generator to a broad range of cache qualities during autoregressive training, we adopt a mixed-step rollout strategy: at each iteration, we sample $K \in \{2, 4, 8\}$ and run chunk-wise autoregressive rollout using the corresponding schedule $\mathcal{T}^{(K)}$. This exposes both the generator and the critic to KV-cache patterns produced by different rollout fidelities, including the degraded caches that arise under low- K settings.

Importantly, mixed-step rollout also provides the training scaffold for applying SC-DMD in the autoregressive setting. We apply the standard DMD updates on generated samples for all sampled step counts, but activate the SC-DMD semigroup-defect regularizer only when $K=8$ is sampled. This design targets the longest composition chain, where multi-step inconsistency is most pronounced and self-consistency regularization is most beneficial. Thus, mixed-step rollout serves two roles simultaneously: it broadens the cache regimes seen during training, and it enables SC-DMD to regularize the high-step regime that matters most for stable autoregressive generation.

Cache-conditioned reference alignment. To further improve robustness, we introduce a *cache-conditioned reference alignment* loss that regularizes outputs under weaker cache conditioning toward those under a stronger reference cache. Given the same input noise ϵ , the generator’s output quality is largely governed by the quality of the conditioning KV cache. For rollouts sampled at $K \in \{2, 4\}$, we construct a higher-quality reference using the next denser schedule, *i.e.* $2 \rightarrow 4$ and $4 \rightarrow 8$, and align the resulting outputs in a TRD-style [45] relational feature space. Formally, let

$$z^{\text{low}}, z^{\text{ref}} \in \mathbb{R}^{F \times S \times D}$$

denote the intermediate features extracted from the lower-quality rollout and the higher-quality reference rollout, respectively, at a sampled noise level t_f , where F is the number of frames, S is the number of spatial tokens per frame, and D is the channel dimension. We first ℓ_2 -normalize the features along the channel dimension and compute, for each frame f , the spatial token relation matrices

$$R_f^{\text{low}} = \bar{z}_f^{\text{low}} (\bar{z}_f^{\text{low}})^\top, \quad R_f^{\text{ref}} = \bar{z}_f^{\text{ref}} (\bar{z}_f^{\text{ref}})^\top, \quad (11)$$

where $\bar{z}$ denotes the normalized feature and $R_f^{(\cdot)} \in \mathbb{R}^{S \times S}$ encodes pairwise token similarities within frame f . We then penalize the discrepancy between the two relation matrices with a margin-relaxed objective:

$$\mathcal{L}_{\text{align}} = \frac{1}{F} \sum_{f=1}^F \frac{1}{S^2} \sum_{i=1}^S \sum_{j=1}^S \left[|R_f^{\text{low}}(i, j) - R_f^{\text{ref}}(i, j)| - \delta \right]_+, \quad (12)$$

where $[\cdot]_+ = \max(\cdot, 0)$ and δ is a margin hyperparameter. This loss encourages the lower-quality rollout to preserve the spatial relational structure of the stronger reference while ignoring small mismatches that are likely uninformative.

Final autoregressive training objective. As shown in Fig. 3, the complete training objective for autoregressive distillation is

$$\min_{\theta} \mathcal{L}_{\text{DMD}}(\theta; \psi) + \lambda_{\text{SC}} \mathcal{L}_{\text{SC}}(\theta) \mathbf{1}[K=8] + \lambda_{\text{align}} \mathcal{L}_{\text{align}}(\theta) \mathbf{1}[K \in \{2, 4\}], \quad (13)$$

where $\mathcal{L}_{\text{SC}}$ is applied at $K=8$ to regularize the longest composition chain, and $\mathcal{L}_{\text{align}}$ is activated for low-step rollouts as a soft constraint toward stronger cache-conditioned references. In practice, the reference alignment is applied stochastically for efficiency; detailed scheduling is deferred to the implementation details.

4 Experimental Results

4.1 Experimental Setup

We evaluate our method in both non-autoregressive and autoregressive settings. For non-autoregressive distillation, we study two regimes: image-to-video (I2V) on Wan 2.1 I2V 14B [35] and text-to-video (T2V) on Wan 2.1 T2V 1.3B, both under extremely low inference budgets. For autoregressive real-time generation, we follow the training and evaluation protocols of recent real-time T2V paradigms and apply our method on top of representative backbones including Self Forcing [13], LongLive [39], and Causal Forcing [48]. More implementation details are provided in the supplementary material.

Baselines. In the non-autoregressive setting, we compare against representative few-step video distillation methods including PCM [36], DMD [42], LightX2V [6], and rCM [47]. Unless otherwise specified, methods are compared under the same inference budget. For LightX2V, we follow the official evaluation protocol and report results from the released model. In the autoregressive setting, we build upon strong real-time baselines including Self Forcing, LongLive, and Causal Forcing, and apply the full **Salt** framework (SC-DMD with cache-conditioned training) without modifying model architectures or inference pipelines.

Benchmarks. For non-autoregressive I2V, we train PCM, DMD, and our method on the same proprietary internal I2V training set and evaluate on VBench-I2V from VBench++ [18], reporting the overall I2V Score and Quality Score together with key dimensions such as Subject/Background Consistency, Motion Smoothness, Dynamic Degree, Imaging Quality, and Temporal Flickering. For non-autoregressive and autoregressive T2V, we report VBench [17] scores on 5-second generation. To assess long-horizon transfer, we additionally evaluate 30-second autoregressive rollouts on VBench-Long from VBench++ [18].

4.2 Non-autoregressive Few-Step Distillation

Image-to-Video on Wan 2.1 14B. Tab. 1 reports VBench-I2V results under an extremely low inference budget of 4 NFEs. Under the same training setting as DMD (8 H200 GPUs, 800 iterations), SC-DMD achieves the best *I2V Score* (93.90) and *Imaging Quality* (72.16), while substantially improving temporal coherence over DMD: *Background Consistency* increases from 92.79 to 95.97, *Motion Smoothness* from 97.99 to 98.37, and *Temporal Flickering* from 95.21 to 97.41. These gains indicate that semigroup-defect regularization improves multi-step compositional stability without sacrificing visual fidelity. Compared with LightX2V, our method is also much more training-efficient, while remaining competitive or superior on most key metrics. Although *Dynamic Degree*

Table 1: VBench-I2V results on Wan 2.1 14B under a 4-NFE inference budget. Higher is better ($\uparrow$).

Method	NFE	VBench-I2V Evaluation Scores $\uparrow$						
		I2V Score	Quality Score	Backg. Consist.	Motion Smooth.	Dynamic Degree	Imaging Quality	Temporal Flicker.
<i>Consistency distillation</i>								
PCM [36]	8	93.63	78.52	97.34	98.24	30.98	70.42	97.67
<i>Distribution matching</i>								
DMD [42]	4	93.09	78.89	92.79	97.99	58.46	70.35	95.21
LightX2V [6]	4	93.50	<u>80.92</u>	95.87	97.89	<u>60.33</u>	71.67	96.30
Ours	4	93.90	80.86	<u>95.97</u>	98.37	52.85	72.16	<u>97.41</u>
Ours-α	4	<u>93.88</u>	81.71	95.46	<u>98.30</u>	68.13	<u>72.08</u>	96.48

decreases relative to DMD (58.46 $\rightarrow$ 52.85), the simultaneous gains in smoothness, flicker suppression, and imaging quality suggest that SC-DMD favors more coherent motion rather than simply increasing optical-flow magnitude.

We also report an adversarial variant, Ours- α , which adds an extra discriminator during training. It further improves *Quality Score* to 81.71 and *Dynamic Degree* to 68.13, showing that SC-DMD is compatible with stronger perceptual objectives. More implementation details are provided in the supplementary material.

Text-to-Video on Wan 2.1 1.3B. We further evaluate SC-DMD in the non-autoregressive T2V setting (“Diffusion models” group in Tab. 2). Under the standard 4-NFE setting, SC-DMD improves over DMD on both *Total Score* (83.36 vs. 82.78) and *Quality Score* (84.76 vs. 84.39), while remaining competitive with rCM on *Semantic Score*. More importantly, the gain persists even at 2 NFEs: SC-DMD improves *Total Score* from 82.41 to 82.85 and *Quality Score* from 83.49 to 84.06 over DMD. These results show that the benefit of compositional regularization remains effective even under more aggressive few-step compression.

4.3 Autoregressive Real-Time Video Generation

5-second short video generation. Tab. 2 reports VBench scores for single-prompt 5-second generation. We apply our method to three representative autoregressive backbones: Self Forcing [13], LongLive [39], and Causal Forcing [48], without changing their architectures or inference pipelines. Under the standard 4-NFE setting, our method consistently improves *Total Score* and *Quality Score* across all three backbones. On LongLive, *Total Score* improves from 84.40 to 84.93 and *Quality Score* from 85.12 to 85.41; on Self Forcing, *Total Score* improves from 84.20 to 84.47 and *Quality Score* from 84.74 to 85.27; on Causal Forcing, *Total Score* improves from 84.62 to 85.08 and *Quality Score* from 85.41 to 85.96. The strongest semantic gain appears on LongLive (81.53 $\rightarrow$ 83.00), while Causal Forcing also improves in *Semantic Score* (81.47 $\rightarrow$ 81.59). As a complementary human-feedback-oriented evaluation, Salt also consistently improves

Table 2: VBench results for 5-second video generation with representative open-source models of similar parameter sizes and resolutions. Higher is better ($\uparrow$).

Model	#Params	Resolution	NFE	VBench Scores $\uparrow$		
				Total	Quality	Semantic
<i>Diffusion models</i>						
rCM [47]	1.3B	832×480	4	82.73	83.65	79.04
DMD [42]	1.3B	832×480	4	82.78	84.39	76.36
Ours (SC-DMD)	1.3B	832×480	2	82.41	83.49	78.06
	1.3B	832×480	4	83.36	84.76	77.77
	1.3B	832×480	2	82.85	84.06	78.01
<i>Autoregressive models</i>						
CausVid [43]	1.3B	832×480	4	81.20	84.05	69.80
Self Forcing [13]	1.3B	832×480	4	84.20	84.74	82.05
Ours - Self Forcing	1.3B	832×480	4	84.47	85.27	81.28
LongLive [39]	1.3B	832×480	4	84.40	85.12	81.53
Ours - LongLive	1.3B	832×480	4	84.93	85.41	83.00
Causal Forcing [48]	1.3B	832×480	4	84.62	85.41	81.47
Ours - Causal Forcing	1.3B	832×480	4	85.08	85.96	81.59
	1.3B	832×480	2	<u>84.80</u>	<u>85.63</u>	<u>81.49</u>

VideoAlign [23] and its motion-quality dimension across the three autoregressive backbones; detailed results are provided in the supplementary material.

The benefit also extends to more aggressive inference budgets. In the Causal Forcing family, our 2-step model reaches a *Total Score* of 84.80 and a *Quality Score* of 85.63, surpassing the original 4-step Causal Forcing baseline (84.62 Total, 85.41 Quality). This shows that our training objective remains effective even when the inference budget is further reduced.

Visualization results. Fig. 4 shows two challenging cases: one dominated by fine-grained textures and the other by high-dynamic motion. In the snowflake example, the LongLive baseline progressively loses the detailed wool texture and produces visible blurring, whereas Ours-LongLive preserves the fine-grained fabric structure. In the train example, the Self Forcing baseline exhibits structural distortion under rapid motion, while Ours-SF maintains a coherent locomotive shape and more stable scene geometry. These examples qualitatively support the gains in visual fidelity and temporal stability reported in Tab. 2.

30-second long video generation. Although our method is designed for real-time autoregressive generation rather than explicit long-video optimization, we evaluate whether its benefits transfer to longer horizons. Following CausVid [43], we select 128 prompts from MovieGenBench and generate five 30-second videos per prompt, with Infinity-RoPE applied to both baselines and our method as a shared long-generation setup. Results on VBench-Long are reported in Tab. 3.

The gains carry over to long-horizon generation. On LongLive, our method improves *Total Score* from 79.03 to 79.27, *Quality Score* from 82.82 to 82.90, and *Semantic Score* from 63.88 to 64.74. On Causal Forcing, *Total Score* improves

Table 3: VBench-Long results for 30-second autoregressive video generation. Higher is better ($\uparrow$). **Bold** indicates the best within each family.

Model	VBench-Long Evaluation Scores $\uparrow$								
	Total Score	Quality Score	Semantic Score	Subject Consist.	Temp. Flicker.	Motion Smooth.	Dynamic Degree	Aesthetic Quality	Imaging Quality
<i>Causal Forcing family</i>									
Causal Forcing [48]	78.11	82.57	60.25	96.29	95.47	97.67	76.95	56.50	71.12
Ours - Causal Forcing	78.28	82.15	62.77	97.16	96.95	98.32	54.02	57.42	70.99
<i>LongLive family</i>									
LongLive [39]	79.03	82.82	63.88	97.36	96.62	98.16	63.45	58.26	70.90
Ours - LongLive	79.27	82.90	64.74	97.28	96.92	98.29	63.68	58.35	70.07




Fig. 4: Qualitative comparison on texture-rich and high-dynamic scenes.

from 78.11 to 78.28, while *Semantic Score* increases substantially from 60.25 to 62.77. These results suggest that the per-chunk compositional gains from SC-DMD and the broader cache exposure from mixed-step training transfer beyond short-horizon rollout, especially in overall quality and semantic consistency.

4.4 Ablation Study

Effect of $\mathcal{L}_{SC}$: fixing the compositionality deficit. We first ablate SC-DMD in a clean non-autoregressive setting on Wan 2.1 1.3B T2V, where all variants use the same 4-step inference budget. Tab. 4 compares three students: DMD-4, trained on a 4-point grid; DMD-8, trained on a denser 8-point grid; and SC-DMD, which uses the same 8-point grid as DMD-8 but adds the shortcut self-consistency loss. Simply increasing the training grid density does not solve the problem: compared with DMD-4, DMD-8 drops from 84.39 to 84.05 in *Quality* and from 82.78 to 82.54 in *Total*, while also degrading several structure-sensitive dimensions such as *Spatial Relation*, *Multi-Objects*, *Object Class*, and *Imaging Quality*. In contrast, SC-DMD consistently improves over both DMD baselines, achieving the best *Total* (83.36), *Quality* (84.76), and *Semantic* (77.77), together

Table 4: Ablation of SC-DMD on Wan 2.1 1.3B T2V under 4-step inference. Increasing the DMD training grid from 4 to 8 timesteps does not improve few-step generation, whereas SC-DMD consistently improves both overall and structure-sensitive metrics. Higher is better ($\uparrow$).

Method	Quality	Semantic	Total	Spatial Rel.	Multi-Objects	Object Class	Imaging Q.
DMD-8	84.05	76.50	82.54	67.13	78.98	93.47	65.46
DMD-4	84.39	76.36	82.78	69.49	79.37	93.23	64.73
SC-DMD (Ours)	84.76	77.77	83.36	71.91	83.89	93.91	67.32




Fig. 5: Cross-step consistency under the same seed and prompt. Left: DMD. Right: SC-DMD. Rows correspond to 2-, 4-, and 8-step sampling, respectively. Vanilla DMD exhibits strong step-dependent variation and structural corruption as the number of steps changes, whereas SC-DMD yields much more consistent outputs across step counts and better preserves object structure.

with clear gains on the fine-grained structural dimensions. This confirms that the gain comes from explicit compositional regularization rather than from training on a denser timestep grid.

Cross-step compositional consistency. Fig. 5 compares DMD and SC-DMD under the same prompt and random seed across 2-, 4-, and 8-step sampling. Vanilla DMD exhibits clear step-dependent inconsistency: the generated content changes noticeably as the step count varies, and higher-step sampling can introduce structural artifacts such as duplicated heads and unstable object identity. In contrast, SC-DMD produces much more consistent outputs across 2/4/8 steps and is markedly less prone to structural corruption. This qualitative evidence directly supports our central claim: SC-DMD improves the *composition* of learned updates, rather than merely optimizing one particular operating point.

Ablation of autoregressive training components. We next ablate the autoregressive training recipe on Causal Forcing under the same VBench-extend prompt setting as the main AR results. Tab. 5 shows three key observations. First, naively adding $\mathcal{L}_{SC}$ to the official baseline hurts performance (*Total*: $84.62 \rightarrow 83.83$), showing that self-consistency alone is not sufficient in the autoregressive setting. Second, once paired with mixed-step rollout, SC-DMD becomes effective and improves the 4-NFE baseline to 85.02 in *Total*, confirming that mixed-step rollout is the training scaffold that makes SC regularization work

Table 5: Ablation of autoregressive training components on Causal Forcing under the VBench-extend prompt setting. The 4-NFE block shows the main component ablation; the 2-NFE block shows low-budget transfer, where reference alignment becomes more helpful. Higher is better ($\uparrow$).

Method	NFE	Quality	Semantic	Total
<i>4-NFE main ablation</i>				
Causal Forcing (official baseline)	4	85.41	81.47	84.62
Causal Forcing + naive $\mathcal{L}_{SC}$	4	84.35	81.77	83.83
Causal Forcing + mixed-step + $\mathcal{L}_{SC}$	4	85.91	81.48	85.02
Ours (full)	4	85.96	81.59	85.08
<i>2-NFE low-budget transfer</i>				
Causal Forcing + mixed-step + $\mathcal{L}_{SC}$	2	85.62	80.65	84.63
Ours (full)	2	85.63	81.49	84.80

in AR generation. Third, the full model achieves the best 4-NFE result (85.08 *Total*, 81.59 *Semantic*) after adding reference alignment. Notably, the gain from reference alignment becomes more pronounced at 2 NFEs: compared with the mixed-step+ $\mathcal{L}_{SC}$ variant, the full model improves *Semantic* from 80.65 to 81.49 and *Total* from 84.63 to 84.80. This is consistent with the intended role of the alignment term, namely guiding weaker low-step rollouts toward stronger high-step references.

5 Conclusion

We presented **Salt**, a training framework for few-step video diffusion distillation that addresses two structural bottlenecks. First, we identified the *compositionality deficit* of DMD: its purely local supervision does not enforce coherent multi-step generation. To remedy this, we proposed **SC-DMD**, which builds on distribution matching with a lightweight semigroup-defect regularizer that encourages endpoint-consistent composition of denoising updates. Second, for autoregressive generation, we introduced *cache-conditioned training* to handle the quality variability of model-generated KV caches: mixed-step rollout exposes the model to diverse cache regimes during training, while reference alignment regularizes weak-cache outputs toward stronger references. Across both non-autoregressive (Wan 2.1) and autoregressive (Self Forcing, Causal Forcing, and LongLive) settings, **Salt** delivers consistent gains in few-step generation quality, cross-step compositional consistency, and long-horizon semantic stability.

Acknowledgement

This work was supported by the Hong Kong Research Grants Council under the Areas of Excellence scheme grant AoE/E-601/22-R and NSFC/RGC Collaborative Research Scheme grant CRS_HKUST603/22.

References

1. Boffi, N., Albergo, M., Vanden-Eijnden, E.: How to build a consistency model: Learning flow maps via self-distillation. *Advances in Neural Information Processing Systems* **38**, 33346–33382 (2026)
2. Boffi, N.M., Albergo, M.S., Vanden-Eijnden, E.: Flow map matching with stochastic interpolants: A mathematical framework for consistency models. *arXiv preprint arXiv:2406.07507* (2024)
3. Cai, X., Wu, Y., Chen, Q., Wu, H., Xiang, L., Wen, H.: Shortcutting pre-trained flow matching diffusion models is almost free lunch. *Advances in Neural Information Processing Systems* **38**, 153066–153094 (2026)
4. 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)
5. Cheng, J., Ma, B., Ren, X., Jin, H.H., Yu, K., Zhang, P., Li, W., Zhou, Y., Zheng, T., Lu, Q.: Phased one-step adversarial equilibrium for video diffusion models. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 40, pp. 3237–3245 (2026)
6. Contributors, L.: Lightx2v: Light video generation inference framework. <https://github.com/ModelTC/lightx2v> (2025)
7. Frans, K., Hafner, D., Levine, S., Abbeel, P.: One step diffusion via shortcut models. In: *International Conference on Learning Representations*. vol. 2025, pp. 34668–34684 (2025)
8. Ge, X., Zhang, X., Xu, T., Zhang, Y., Zhang, X., Wang, Y., Zhang, J.: Sense-flow: Scaling distribution matching for flow-based text-to-image distillation. *arXiv preprint arXiv:2506.00523* (2025)
9. Geng, Z., Deng, M., Bai, X., Kolter, Z., He, K.: Mean flows for one-step generative modeling. *Advances in Neural Information Processing Systems* **38**, 75460–75482 (2026)
10. Hastie, T., Tibshirani, R., Friedman, J.H., Friedman, J.H.: *The elements of statistical learning: data mining, inference, and prediction*, vol. 2. Springer (2009)
11. 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)
12. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
13. Huang, X., Li, Z., He, G., Zhou, M., Shechtman, E.: Self forcing: Bridging the train-test gap in autoregressive video diffusion. *Advances in Neural Information Processing Systems* **38**, 167283–167308 (2026)
14. Huang, Y., Ge, X., Gong, R., Lv, C., Zhang, J.: Linvideo: A post-training framework towards o(n) attention in efficient video generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 23398–23408 (2026)
15. Huang, Y., Gong, R., Liu, J., Ding, Y., Lv, C., Qin, H., Zhang, J.: Qvgen: Pushing the limit of quantized video generative models. *arXiv preprint arXiv:2505.11497* (2025)
16. Huang, Y., Wang, Z., Gong, R., Liu, J., Zhang, X., Guo, J., Liu, X., Zhang, J.: Harmonica: Harmonizing training and inference for better feature caching in diffusion transformer acceleration. *arXiv preprint arXiv:2410.01723* (2024)

17. Huang, Z., He, Y., Yu, J., Zhang, F., Si, C., Jiang, Y., Zhang, Y., Wu, T., Jin, Q., Chanpaisit, N., Wang, Y., Chen, X., Wang, L., Lin, D., Qiao, Y., Liu, Z.: VBench: Comprehensive benchmark suite for video generative models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2024)
18. Huang, Z., Zhang, F., Xu, X., He, Y., Yu, J., Dong, Z., Ma, Q., Chanpaisit, N., Si, C., Jiang, Y., et al.: Vbench++: Comprehensive and versatile benchmark suite for video generative models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025)
19. Kim, D., Lai, C.H., Liao, W., Murata, N., Takida, Y., Uesaka, T., He, Y., Mitsu-fuji, Y., Ermon, S.: Consistency trajectory models: Learning probability flow ode trajectory of diffusion. In: International Conference on Learning Representations. vol. 2024, pp. 44493–44525 (2024)
20. Lin, S., Xia, X., Ren, Y., Yang, C., Xiao, X., Jiang, L.: Diffusion adversarial post-training for one-step video generation. *arXiv preprint arXiv:2501.08316* (2025)
21. Lin, S., Yang, C., He, H., Jiang, J., Ren, Y., Xia, X., Zhao, Y., Xiao, X., Jiang, L.: Autoregressive adversarial post-training for real-time interactive video generation. *arXiv preprint arXiv:2506.09350* (2025)
22. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: ICLR. OpenReview.net (2023)
23. Liu, J., Liu, G., Liang, J., Yuan, Z., Liu, X., Zheng, M., Wu, X., Wang, Q., Xia, M., Wang, X., et al.: Improving video generation with human feedback. *Advances in Neural Information Processing Systems* **38**, 82155–82192 (2026)
24. Liu, K., Hu, W., Xu, J., Shan, Y., Lu, S.: Rolling forcing: Autoregressive long video diffusion in real time. *arXiv preprint arXiv:2509.25161* (2025)
25. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003* (2022)
26. Liu, Y., Liu, B., Zhang, Y., Hou, X., Song, G., Liu, Y., You, H.: See further when clear: Curriculum consistency model. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 18103–18112 (2025)
27. 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)
28. Luo, S., Tan, Y., Huang, L., Li, J., Zhao, H.: Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378* (2023)
29. Lv, Z., Si, C., Pan, T., Chen, Z., Wong, K.Y.K., Qiao, Y., Liu, Z.: Dual-expert consistency model for efficient and high-quality video generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 14983–14993 (October 2025)
30. Mao, X., Jiang, Z., Wang, F.Y., Zhang, J., Chen, H., Chi, M., Wang, Y., Luo, W.: Osv: One step is enough for high-quality image to video generation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 12585–12594 (2025)
31. Nie, W., Berner, J., Ma, N., Liu, C., Xie, S., Vahdat, A.: Transition matching distillation for fast video generation. *arXiv preprint arXiv:2601.09881* (2026)
32. Sabour, A., Fidler, S., Kreis, K.: Align your flow: Scaling continuous-time flow map distillation. *Advances in Neural Information Processing Systems* **38**, 146459–146512 (2026)
33. Song, Y., Dhariwal, P., Chen, M., Sutskever, I.: Consistency models. In: International Conference on Machine Learning. pp. 32211–32252. PMLR (2023)

34. Teng, H., Jia, H., Sun, L., Li, L., Li, M., Tang, M., Han, S., Zhang, T., Zhang, W., Luo, W., et al.: Magi-1: Autoregressive video generation at scale. arXiv preprint arXiv:2505.13211 (2025)
35. 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)
36. Wang, F.Y., Huang, Z., Bergman, A., Shen, D., Gao, P., Lingelbach, M., Sun, K., Bian, W., Song, G., Liu, Y., et al.: Phased consistency models. *Advances in neural information processing systems* **37**, 83951–84009 (2024)
37. Wang, Y., Zhang, H., Xue, T., Qiao, Y., Wang, Y., Xu, C., Chen, X.: Vdot: Efficient unified video creation via optimal transport distillation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9273–9283 (2026)
38. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Advances in Neural Information Processing Systems* **36**, 8406–8441 (2023)
39. 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. arXiv preprint arXiv:2509.22622 (2025)
40. Yesiltepe, H., Meral, T., Akan, A.K., Oktay, K., Yanardag, P.: Infinity-rope: Action-controllable infinite video generation emerges from autoregressive self-rollout. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 40256–40265 (2026)
41. Yin, T., Gharbi, M., Park, T., Zhang, R., Shechtman, E., Durand, F., Freeman, B.: Improved distribution matching distillation for fast image synthesis. *Advances in neural information processing systems* **37**, 47455–47487 (2024)
42. Yin, T., Gharbi, M., Zhang, R., Shechtman, E., Durand, F., Freeman, W.T., Park, T.: One-step diffusion with distribution matching distillation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 6613–6623 (2024)
43. 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)
44. Zhai, Y., Lin, K., Yang, Z., Li, L., Wang, J., Lin, C.C., Doermann, D., Yuan, J., Wang, L.: Motion consistency model: Accelerating video diffusion with disentangled motion-appearance distillation. *Advances in Neural Information Processing Systems* **37**, 111000–111021 (2024)
45. Zhang, X., Liao, J., Zhang, S., Meng, F., Wan, X., Yan, J., Cheng, Y.: Videorepa: Learning physics for video generation through relational alignment with foundation models. *Advances in Neural Information Processing Systems* **38**, 122647–122676 (2026)
46. Zhao, M., Zhu, H., Wang, Y., Yan, B., Zhang, J., He, G., Yang, L., Li, C., Zhu, J.: Ultravico: Breaking extrapolation limits in video diffusion transformers. arXiv preprint arXiv:2511.20123 (2025)
47. Zheng, K., Wang, Y., Ma, Q., Chen, H., Zhang, J., Balaji, Y., Chen, J., Liu, M.Y., Zhu, J., Zhang, Q.: Large scale diffusion distillation via score-regularized continuous-time consistency. arXiv preprint arXiv:2510.08431 (2025)
48. 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)