

SwiftVideo: A Unified Framework for Few-Step Video Generation Through Trajectory-Distribution Alignment

Yanxiao Sun^{1*}, Jiafu Wu^{2*}, Yun Cao², Chengming Xu², Yabiao Wang², Weijian Cao²
Donghao Luo², Chengjie Wang², Yanwei Fu^{1,3 †}

¹Fudan University

²Tencent Youtu Lab

³Shanghai Innovation Institute

Abstract

Diffusion-based or flow-based models have achieved significant progress in video synthesis but require multiple iterative sampling steps, which incurs substantial computational overhead. While many distillation methods that are solely based on trajectory-preserving or distribution-matching have been developed to accelerate video generation models, these approaches often suffer from performance breakdown or increased artifacts under few-step settings. To address these limitations, we propose *SwiftVideo*, a unified and stable distillation framework that combines the advantages of trajectory-preserving and distribution-matching strategies. Our approach introduces continuous-time consistency distillation to ensure precise preservation of ODE trajectories. Subsequently, we propose a dual-perspective alignment that includes distribution alignment between synthetic and real data along with trajectory alignment across different inference steps. Our method maintains high-quality video generation while substantially reducing the number of inference steps. Quantitative evaluations on the OpenVid-1M benchmark demonstrate that our method significantly outperforms existing approaches in few-step video generation.

Introduction

Recent advancements in diffusion and flow-based models have significantly enhanced generative modeling, especially in video synthesis. However, these models typically require many iterative sampling steps, resulting in high computational and time costs that limit their practical deployment. Although higher-order numerical ODE solvers (Lu et al. 2025, 2022) can reduce the number of sampling steps, they often suffer from severe quality degradation when operating in the few-step regime (i.e., steps ≤ 10).

To address these challenges, diffusion distillation methods have emerged to reduce sampling overhead in video generation. Trajectory-preserving approaches, such as LCM, PCM (Luo et al. 2023; Wang et al. 2024), leverage consistency loss to minimize inference steps. However, as discrete-time consistency models (CMs), they tend to produce blurry

*These authors contributed equally.

†Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

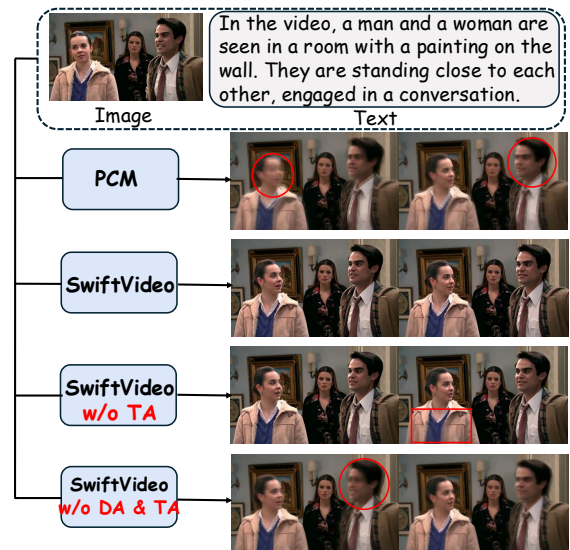


Figure 1: Video generation results on the OpenVid-1M dataset (Nan et al. 2025) using 4 sampling steps. Red circles and boxes highlight blurry regions and artifacts in the generated videos. DA and TA denote Distribution Alignment and Trajectory Alignment, respectively. SwiftVideo w/o DA & TA employs only Continuous-time Consistency Distillation (CCD). Our method significantly outperforms PCM.

outputs under few-step settings and introduce additional discretization errors from numerical ODE solvers. Alternatively, distribution-matching methods like DMD2 (Yin et al. 2024a), train student models to approximate the teacher’s distribution in order to reduce inference steps. However, this approach inherently upper-bounds the student model’s performance by the teacher model and can lead to domain inconsistencies due to relying solely on distributional constraints. Other recent strategies (Ren et al. 2024; Sauer et al. 2024a) incorporate post-training enhancements for few-step generation, but often depend on auxiliary reward models or expensive-to-collect preference datasets.

Based on this analysis, we identify three key limitations in existing methods: 1) Discretization errors from numerical ODE solvers in previous CMs; 2) Performance ceilings

imposed by teacher model distributions; 3) Quality degradation under extremely few steps. To this end, we propose *SwiftVideo*, a unified and stable distillation framework that systematically addresses these limitations. Our framework comprises three components: continuous-time consistency distillation (CCD) for precise ODE trajectory preservation, distribution alignment (DA) on real data to enhance visual fidelity, and trajectory alignment (TA) across different inference steps for enhanced few-step generation.

Firstly, we introduce CCD to ensure precise ODE trajectory preservation, eliminating the discretization errors that arise from ODE solver reliance in discrete-time consistency models. As shown in Figure 1, while CCD alone outperforms PCM, it still exhibits blurriness. Nevertheless, its accurate trajectory preservation enables effective subsequent training optimization. Secondly, instead of approximating the teacher’s distribution, we directly perform DA on real data, thereby breaking the performance constraints imposed by the teacher model. Thirdly, we observe that the distribution-aligned distilled model exhibits significant improvement in generation quality as the number of steps increases within the low-step regime (i.e., steps ≤ 8). Based on this observation, we propose TA across different inference steps as a post-training strategy. Specifically, we generate the preference dataset using the distilled model with varying numbers of steps, then leverage the Direct Preference Optimization (DPO) algorithm (Rafailov et al. 2023) to implicitly align low-step inference trajectories with high-step inference trajectories, eliminating the need for reward models and costly external preference data collection. As shown in Figure 1, DA effectively reduces blurriness while TA improves detail preservation under few-step settings.

In summary, our main contributions are as follows: (1) We propose a unified and stable distillation framework that maintains high-quality video generation while significantly reducing the number of inference steps. (2) **Continuous-time Consistency Distillation (CCD)**: We effectively apply CCD to video generation models for the first time, ensuring precise ODE trajectory preservation while incorporating stabilization and acceleration techniques throughout the training process. (3) **Distribution Alignment (DA)**: We employ adversarial training to enable the model to directly align with real data distributions, thereby enhancing the visual details and realism of generated results. (4) **Trajectory Alignment (TA)**: We apply the DPO algorithm on synthetic preference dataset across different inference steps to enforce alignment with high-step inference trajectories, thus enhancing few-step generation quality. (5) Extensive experiments demonstrate that our approach outperforms existing methods in few-step video generation, achieving state-of-the-art results on VBench and FVD benchmarks.

Related Work

Video Diffusion Model

Driven by advances in large-scale video data and model architectures, video diffusion models have achieved remarkable success. Early approaches (Blattmann et al. 2023; Guo et al. 2024) using 3D U-Net architectures in latent

space were limited to generating short, low-quality videos. With the success of Sora (OpenAI 2024), modern models have shifted toward Diffusion Transformer (DiT) architectures, with open-source implementations like Hunyuan-Video (Kong et al. 2025), CogVideoX (Yang et al. 2025), and Wan (Wan et al. 2025). Although these models can produce longer, higher-quality videos, generation remains computationally expensive and time-consuming. To address this limitation, we leverage diffusion distillation methods to accelerate these models without compromising their quality.

Diffusion Model Distillation

Diffusion models typically rely on iterative inference to generate high-quality samples, which is both time-consuming and computationally expensive. Diffusion model distillation has been regarded as an effective technique for accelerating the inference process and can be categorized into trajectory-preserving and distribution-matching approaches.

Trajectory-preserving distillation aims to train a student model capable of generating samples in fewer steps by approximating the ordinary differential equation (ODE) sampling trajectory of the teacher model. Progressive Distillation (Salimans and Ho 2022) trains the student model to predict the subsequent flow locations determined by the teacher model. Rectified Flow (Liu, Gong, and Liu 2022) trains the student model on noisy-image pairs obtained from the teacher, thereby progressively straightening the ODE trajectory. Consistency Distillation (Luo et al. 2023) trains a student model to map noisy samples along the ODE trajectory to its origin. Although trajectory-preserving distillation facilitates training, the samples generated in few-step tend to exhibit blurriness and inferior quality.

Distribution-matching distillation trains a student model to approximate the teacher’s distribution. Some methods leverage adversarial loss functions to improve the generative quality of student models (Sauer et al. 2024b,a). DMD (Yin et al. 2024b) minimizes approximate KL divergence using gradients from two score functions: a fixed pretrained target score and a dynamically trained generator score.

Recently, trajectory-preserving distillation methods (Li et al. 2024; Wang et al. 2023) and distribution-matching distillation (Zhang et al. 2024) have been increasingly applied to video generative models. However, many of these distilled video models originate from the pre-Sora era and typically generate only low-quality, short-duration results. Moreover, existing consistency distillation methods for video are trained using discretized timesteps, which rely on numerical ODE solvers and introduce additional discretization errors. In contrast, our approach employs continuous-time consistency distillation to address these issues. Furthermore, prior distribution-matching distillation methods usually approximate the teacher’s output distribution as the target, which implicitly limits the generation quality by the teacher model. Instead, our approach directly performs distribution alignment on real data, similar to (Lin et al. 2025).

Preliminary

Flow Matching Flow Matching (Lipman et al. 2023; Liu, Gong, and Liu 2022) is widely used in recent advanced

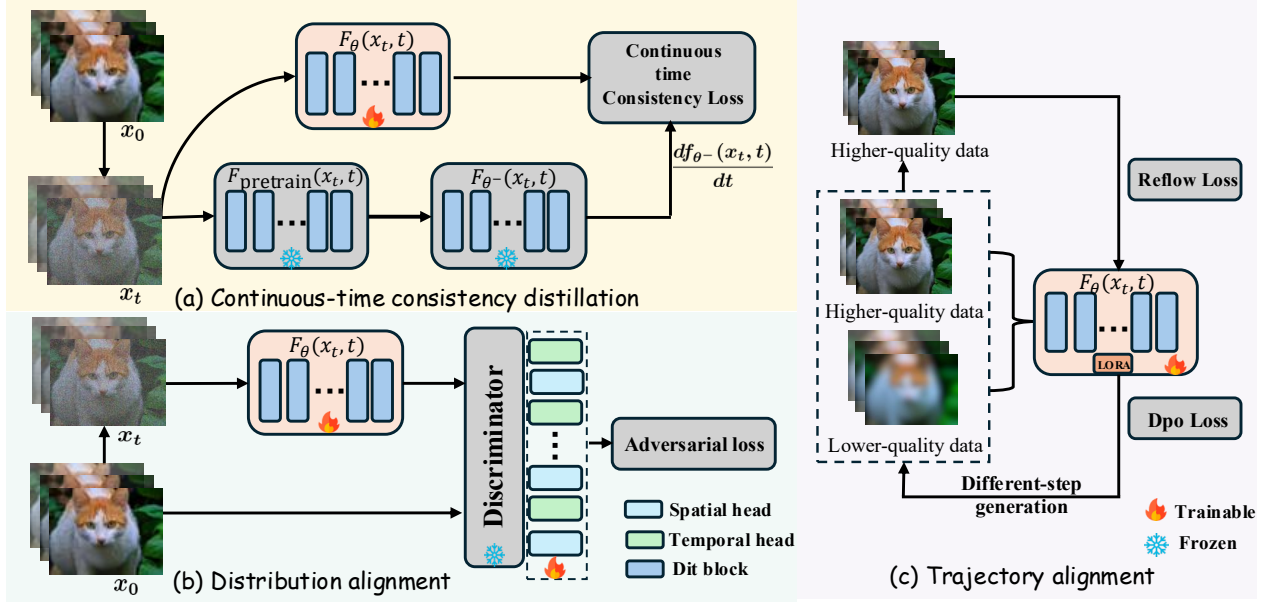


Figure 2: The overview of our distillation framework SwiftVideo.

video-generation models (Wan et al. 2025). Given a data sample $\mathbf{x}_0 \sim p_{\text{data}}$ and a noise sample $\mathbf{x}_1 \sim p_{\text{noise}}$, the noisy data \mathbf{x}_t is defined as a linear interpolation between these two distributions:

$$\mathbf{x}_t = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1, \quad (1)$$

for $t \in [0, 1]$. The target velocity field is $\mathbf{v} = \mathbf{x}_1 - \mathbf{x}_0$. The optimization process of flow matching is defined as:

$$\mathcal{L}(\theta) = \mathbb{E}_{t, \mathbf{x}_0 \sim X_0, \mathbf{x}_1 \sim X_1} [\|\mathbf{v} - \mathbf{v}_\theta(\mathbf{x}_t, t)\|^2], \quad (2)$$

where \mathbf{v}_θ denotes the model parameterized by θ .

Consistency Models A consistency model (CM) is a neural network $\mathbf{f}_\theta(\mathbf{x}_t, t)$ trained to map any points on a trajectory of the PF-ODE to the trajectory’s origin. The formula can be described as $\mathbf{f} : (\mathbf{x}_t, t) \mapsto \mathbf{x}_0$. A valid \mathbf{f}_θ must satisfy the boundary condition $\mathbf{f}_\theta(\mathbf{x}, 0) \equiv \mathbf{x}$, which can be achieved through the parameterization $\mathbf{f}_\theta(\mathbf{x}_t, t) = c_{\text{skip}}(t)\mathbf{x}_t + c_{\text{out}}(t)\mathbf{F}_\theta(c_{\text{in}}(t)\mathbf{x}_t, c_{\text{noise}}(t))$, where $c_{\text{skip}}(0) = 1$ and $c_{\text{out}}(0) = 0$. Depending on the selection of nearby time steps, consistency models can be classified into two categories, as described below:

Discrete-time CMs The training objective defined on two adjacent time steps with a finite distance is as follows:

$$\mathbb{E}_{\mathbf{x}_t, t} [w(t)d(\mathbf{f}_\theta(\mathbf{x}_t, t), \mathbf{f}_\theta^-(\mathbf{x}_{t-\Delta t}, t-\Delta t))] \quad (3)$$

where θ^- denotes $\text{stopgrad}(\theta)$. The term $w(t)$ is a weighting function, $\Delta t > 0$ is the distance between adjacent timesteps, and $d(\cdot, \cdot)$ is a distance function, e.g., the squared ℓ_2 distance $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2^2$. The noisy sample $\mathbf{x}_{t-\Delta t}$ at the preceding time step $t - \Delta t$ is typically obtained by solving the PF-ODE using numerical ODE solvers with a step size of Δt , which introduce additional discretization error.

Continuous-time CMs When taking the limit $\Delta t \rightarrow 0$, show that the gradient of Eq (3) with respect to θ converges to

$$\nabla_{\theta} E_{\mathbf{x}_t, t} \left[w(t) \mathbf{f}_\theta^\top(\mathbf{x}_t, t) \frac{d\mathbf{f}_\theta^-(\mathbf{x}_t, t)}{dt} \right], \quad (4)$$

where $\frac{d\mathbf{f}_\theta^-(\mathbf{x}_t, t)}{dt} = \nabla_{\mathbf{x}_t} \mathbf{f}_\theta^-(\mathbf{x}_t, t) \frac{d\mathbf{x}_t}{dt} + \partial_t \mathbf{f}_\theta^-(\mathbf{x}_t, t)$ is the tangent of \mathbf{f}_θ^- at (\mathbf{x}_t, t) along the trajectory of the PF-ODE $\frac{d\mathbf{x}_t}{dt}$. Continuous-time CMs do not rely on ODE solvers, which avoids discretization errors and offers more accurate supervision signals during training. Recently, sCM (Lu and Song 2025) presents several training tricks to simplify, stabilize, and scale the training of continuous-time consistency models, achieving superior performance compared to discrete-time consistency models.

Method

As shown in Figure 2, our framework systematically addresses the three limitations identified above through the following components: continuous-time consistency distillation, distribution alignment via adversarial training, and DPO-based trajectory alignment with reflow loss regularization.

Continuous-time Consistency Distillation

In previous work, the consistency model \mathbf{f}_θ is parameterized as:

$$\mathbf{f}_\theta(\mathbf{x}_t, t) = c_{\text{skip}}(t)\mathbf{x}_t + c_{\text{out}}(t)\mathbf{F}_\theta(\mathbf{x}_t, t) \quad (5)$$

$\mathbf{F}_\theta(x, t)$ is a deep neural network. Unlike the TrigFlow formulation employed in sCM (Lu and Song 2025), we set $c_{\text{skip}} = 1$ and $c_{\text{out}} = -t$, which aligns with the Rectified Flow (Liu, Gong, and Liu 2022) formulation commonly adopted in current video diffusion models and eliminates

Algorithm 1: Continuous-time Consistency Distillation

1: **Input:** Training dataset D , pretrained model $\mathbf{F}_{\text{pretrain}}$, distilled model \mathbf{F}_θ , learning rate η , timestep schedule $(P_{\text{mean}}, P_{\text{std}}^2)$, consistency loss warmup iteration H
2: **Initialize** $\text{Iters} \leftarrow 0, \theta \leftarrow \phi$
3: **repeat**
4: $\mathbf{x}_0 \sim D, \mathbf{x}_1 \sim \mathcal{N}(0, \mathbf{I}), t \sim \mathcal{N}(P_{\text{mean}}, P_{\text{std}}^2)$
5: $\mathbf{x}_t \leftarrow (1-t)\mathbf{x}_0 + t\mathbf{x}_1, \frac{d\mathbf{x}_t}{dt} \leftarrow \mathbf{F}_{\text{pretrain}}(\mathbf{x}_t, t)$
6: $r \leftarrow \min(1, \text{Iters}/H)$
7: $\mathbf{g} \leftarrow \frac{d\mathbf{x}_t}{dt} - \mathbf{F}_{\theta^-}(\mathbf{x}_t, t) - rt \cdot \text{sg}\left(\frac{d\mathbf{F}_{\theta^-}(\mathbf{x}_t, t)}{dt}\right)$
8: $\mathbf{g} \leftarrow \mathbf{g}/(\|\mathbf{g}\| + c)$
9: $\mathcal{L}_{\text{distill}} \leftarrow \|\mathbf{F}_\theta(\mathbf{x}_t, t) - \mathbf{F}_{\theta^-}(\mathbf{x}_t, t) - \mathbf{g}\|_2^2$
10: $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}_{\text{distill}}$
11: $\text{Iters} \leftarrow \text{Iters} + 1$
12: **until** convergence

the need for additional formulation conversion. Moreover, this parameterization satisfies the boundary conditions and is compatible with the Euler ODE solver. Then we differentiate both sides of Eq (4) with respect to t :

$$\frac{d\mathbf{f}_{\theta^-}(\mathbf{x}_t, t)}{dt} = \frac{d\mathbf{x}_t}{dt} - \mathbf{F}_{\theta^-}(\mathbf{x}_t, t) - t \frac{d\mathbf{F}_{\theta^-}(\mathbf{x}_t, t)}{dt} \quad (6)$$

In this work, we aim to distill a pretrained video generation model into a few-step generator, so $\frac{d\mathbf{x}_t}{dt} = \mathbf{F}_{\text{pretrain}}(\mathbf{x}_t, t)$ in Eq. (6).

Given the equation $\nabla_\theta \mathbb{E}[\mathbf{F}_\theta^\top \mathbf{y}] = \frac{1}{2} \nabla_\theta \mathbb{E}[\|\mathbf{F}_\theta - \mathbf{F}_{\theta^-} + \mathbf{y}\|_2^2]$, where θ^- denotes *stopgrad*(θ) and \mathbf{y} is an arbitrary vector independent of θ . Then we can derive the training objective for continuous-time consistency distillation:

$$\begin{aligned} & \nabla_\theta \mathbb{E}_{\mathbf{x}_t, t} \left[\mathbf{f}_\theta^\top(\mathbf{x}_t, t) \frac{d\mathbf{f}_{\theta^-}(\mathbf{x}_t, t)}{dt} \right] \\ &= -t \nabla_\theta \mathbb{E}_{\mathbf{x}_t, t} \left[\mathbf{F}_\theta^\top(\mathbf{x}_t, t) \frac{d\mathbf{f}_{\theta^-}(\mathbf{x}_t, t)}{dt} \right] \\ &= \frac{t}{2} \nabla_\theta \mathbb{E}_{\mathbf{x}_t, t} \|\mathbf{F}_\theta(\mathbf{x}_t, t) - \mathbf{F}_{\theta^-}(\mathbf{x}_t, t) - \frac{d\mathbf{f}_{\theta^-}(\mathbf{x}_t, t)}{dt}\|_2^2 \end{aligned} \quad (7)$$

However, directly training on this objective suffers from instability and inefficiency issues, it's necessary to incorporate additional training techniques for stabilization and acceleration.

Training Stabilization and Acceleration Building upon the approach proposed in sCM (Lu and Song 2025), we employed two training techniques, Tangent Normalization and Tangent Warmup, to mitigate the instability issues in the continuous-time consistency distillation process. Additionally, given that current video diffusion models comprise multiple DiT blocks, we designed an iterative algorithm to compute the tangent function $\frac{d\mathbf{f}_{\theta^-}(\mathbf{x}_t, t)}{dt}$ with minimal computational overhead.

In our experiments, we found that the term $\frac{d\mathbf{F}_{\theta^-}(\mathbf{x}_t, t)}{dt}$ within the tangent function $\frac{d\mathbf{f}_{\theta^-}(\mathbf{x}_t, t)}{dt}$ is the primary cause of instability in the training objective of continuous-time

Flow Distillation. This term can be derived as follows using the chain rule:

$$\frac{d\mathbf{F}_{\theta^-}(\mathbf{x}_t, t)}{dt} = \nabla_{\mathbf{x}_t} \mathbf{F}_{\theta^-}(\mathbf{x}_t, t) \frac{d\mathbf{x}_t}{dt} + \partial_t \mathbf{F}_{\theta^-}(\mathbf{x}_t, t) \quad (8)$$

Firstly, we employed the Tangent Warmup training trick by introducing a coefficient r to scale the term $t \frac{d\mathbf{F}_{\theta^-}(\mathbf{x}_t, t)}{dt}$ in the tangent function Eq.(6). This coefficient r linearly increases from 0 to 1 over the first 1K training steps. Thus, the tangent function is transformed into the following form.

$$\begin{aligned} & \frac{d\mathbf{x}_t}{dt} - \mathbf{F}_{\theta^-} - rt \frac{d\mathbf{F}_{\theta^-}(\mathbf{x}_t, t)}{dt} \\ &= (1-r) \left(\frac{d\mathbf{x}_t}{dt} - \mathbf{F}_{\theta^-} \right) + r \left(\frac{d\mathbf{x}_t}{dt} - \mathbf{F}_{\theta^-} - t \frac{d\mathbf{F}_{\theta^-}(\mathbf{x}_t, t)}{dt} \right) \end{aligned} \quad (9)$$

We observe that Eq.(9) is expressed as a weighted sum of flow matching loss and original training objective of continuous-time consistency distillation in Eq.(7).

Secondly, we employed the Tangent Normalization to reduce the variance during training, experimenting with two approaches: clipping the tangent within the range of -1 to 1 versus replacing $\frac{d\mathbf{f}_{\theta^-}(\mathbf{x}_t, t)}{dt}$ with $\frac{d\mathbf{f}_{\theta^-}(\mathbf{x}_t, t)}{dt} / (\|\frac{d\mathbf{f}_{\theta^-}(\mathbf{x}_t, t)}{dt}\| + c)$, $c = 0.1$. We found that the latter approach yielded better results.

While the above two training tricks stabilize the training process, the computation of $\frac{d\mathbf{F}_{\theta^-}(\mathbf{x}_t, t)}{dt}$ that lead to efficiency and memory issues when directly applied to video diffusion models. To address this, we apply stop-gradient operations in JVP computation and perform iterative block-wise JVP calculation across DiT blocks, using the JVP output from each block as input for the subsequent block's JVP computation while detaching intermediate results to manage memory consumption. The detailed process for JVP computation is described in the appendix.

Algorithm 1 presents our continuous-time consistency distillation algorithm.

Alignment for Distilled Model

To further improve generation quality in few-step sampling, we align the distilled model from both distribution and trajectory perspectives.

Distribution Alignment While continuous-time consistency distillation achieves precise ODE trajectory preservation, the generated results still exhibit blurriness due to the inherent properties of consistency models (Lin, Wang, and Yang 2024). Therefore, we introduce distribution alignment to improve the visual fidelity and realism of generated videos. Unlike previous distribution-matching methods that constrain the upper limit of the student's generation quality by matching the teacher's distribution, we directly employ adversarial loss to align with real data distribution.

The distribution alignment employs a discriminator trained to distinguish between generated samples $\hat{\mathbf{x}}_0$ and real data \mathbf{x}_0 , where $\hat{\mathbf{x}}_0 = \mathbf{x}_t - t\mathbf{F}_\theta(\mathbf{x}_t, t)$. For the discriminator, we follow the design in StyleGAN-T (Sauer et al. 2023). We utilize a frozen pretrained feature network DINOv2 (Oquab

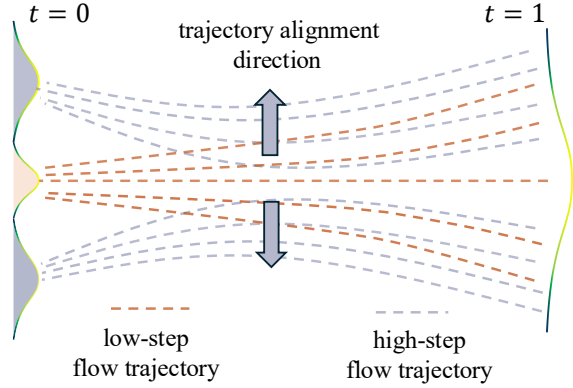


Figure 3: Illustration of trajectory alignment.

et al. 2024) coupled with trainable lightweight discriminator heads, comprising both temporal heads for capturing inter-frame temporal dynamics and spatial heads for evaluating intra-frame visual quality.

For training stability, we introduce distribution alignment after $\mathcal{N}_{\text{warmup}}$ training steps of continuous-time consistency distillation. The objective of the distribution alignment training is as follows:

$$\begin{aligned} \mathcal{L}^{\mathcal{G}} &= \lambda_{\text{adv}} \cdot \mathbb{E}_{\mathbf{x}_0, t} [\sum_k \max(0, 1 - \mathcal{D}_{\phi, k}(\hat{\mathbf{x}}_0))] + \mathcal{L}_{\text{distill}} \\ \mathcal{L}^{\mathcal{D}} &= \mathbb{E}_{\mathbf{x}_0, t} [\sum_k \max(0, 1 - \mathcal{D}_{\phi, k}(\mathbf{x}_0)) \\ &\quad , \max(0, 1 + \mathcal{D}_{\phi, k}(\hat{\mathbf{x}}_0))] \end{aligned} \quad (10)$$

where λ_{adv} is a hyper-parameter, and $\mathcal{D}_{\phi, k}$ denotes either spatial or temporal discriminator heads.

By directly aligning distributions through adversarial training on real data, our approach significantly enhances visual fidelity and realism in video generation.

Trajectory Alignment While distribution alignment enhances the generation quality of distilled models under limited inference steps, visual artifacts persist at extremely low step counts (i.e., steps ≤ 4). We observe that distilled model using Euler solver inference exhibits significant quality improvements as the number of steps increases within the low-step regime (i.e., steps ≤ 8). Based on this observation, we introduce trajectory alignment across different inference steps to advance model’s low-step generation frontier. Initially, we utilize the distribution-aligned distilled model to efficiently construct a synthetic preference dataset $\mathcal{D} = \{\mathbf{y}, \mathbf{x}_0^w, \mathbf{x}_0^l\}$ across varying inference steps. In this preference dataset, each sample contains a prompt \mathbf{y} and two videos \mathbf{x}_0^w and \mathbf{x}_0^l generated by the distribution-aligned distilled model using different steps (e.g., 8 steps and 4 steps) respectively.

Subsequently, we finetune our model using Direct Preference Optimization (Liu et al. 2025) on this synthetic preference dataset. As shown in Figure 3, since the synthetic data implicitly contains flow trajectory information, this training enables the model to align with high-step inference trajectories, thereby enhancing low-step generation quality. The DPO loss \mathcal{L}_{DPO} on flow-based video generation model is

written below:

$$\begin{aligned} \mathcal{L}_{\text{DPO}} &= -\mathbb{E} \left[\log \sigma \left(-\frac{\beta}{2} \right. \right. \\ &\quad \left. \left. \left(\|\mathbf{v}^w - \mathbf{v}_{\theta}(\mathbf{x}_t^w, t)\|^2 - \|\mathbf{v}^w - \mathbf{v}_{\text{ref}}(\mathbf{x}_t^w, t)\|^2 \right. \right. \right. \\ &\quad \left. \left. \left. - \left(\|\mathbf{v}^l - \mathbf{v}_{\theta}(\mathbf{x}_t^l, t)\|^2 - \|\mathbf{v}^l - \mathbf{v}_{\text{ref}}(\mathbf{x}_t^l, t)\|^2 \right) \right) \right) \right], \end{aligned} \quad (11)$$

where β is a hyperparameter, $\{\mathbf{x}_0^w, \mathbf{x}_0^l\} \sim \mathcal{D}$ and $\mathbf{x}_t^* = (1 - t)\mathbf{x}_0^* + t\epsilon$, $\epsilon \sim \mathcal{N}(0, \mathbf{I})$, with $*$ denoting either w (for the preferred data) or l (for the less preferred data).

Employing Eq. (11) alone as the training objective for trajectory alignment results in undesired training dynamics. As training progresses, the model minimizes the loss by reducing velocity prediction accuracy on dispreferred data more than on preferred data, thereby causing the flow to diverge from the intended trajectory. To address this issue, we introduce a regularization term $\|\mathbf{v}^w - \mathbf{v}_{\theta}(\mathbf{x}_t^w, t)\|^2$ to maintain high prediction fidelity on the preferred data. Given that our preference dataset is entirely synthetic, this regularization term takes the same form as the reflow loss (Liu, Gong, and Liu 2022), denoted as \mathcal{L}_{RF} . Therefore, the overall trajectory alignment loss can be formulated as $\mathcal{L}_{\text{DPO}} + \lambda_{\text{RF}} \cdot \mathcal{L}_{\text{RF}}$, where λ_{RF} is a hyper-parameter.

Experiment

Implementation Details. In our distillation experiments, we use Wan2.1-FUN-inp-480p-1.3B (Wan et al. 2025) as the foundational model. For fair comparison, all the experiments are conducted on OpenVid-1M (Nan et al. 2025), a large-scale, high-quality video dataset that provides a standardized benchmark for few-step video generation. Specifically, we randomly select one million videos for training and 1000 for evaluation. In the continuous-time consistency distillation and distribution alignment, we fix the resolution of the training videos to 832×480 with 61 frames and a batch size of 1. We conduct full-parameter training for 3,000 iterations using a learning rate of $1e - 6$ for both the discriminator and student model. The default warm-up iteration $\mathcal{N}_{\text{warmup}}$ is 1000 and λ_{adv} is set to 0.01. In trajectory alignment, we first utilize distribution-aligned distilled model to generate videos using 4 steps and 8 steps sampling respectively, then construct preference dataset of size 5000. Then we apply LoRA (Hu et al. 2022) to fine-tune the model’s linear layers on this preference dataset, using a learning rate of $1e - 6$ for approximately 2000 iterations. To further improve 2-step generation quality, we conduct an additional round of trajectory alignment, employing the once-aligned model to generate preference data using 2-step and 4-step sampling with the same training procedure. λ_{RF} is set to 2 and β in \mathcal{L}_{DPO} is set to 2500. All training uses the AdamW optimizer (Loshchilov and Hutter 2019), and we employ the Euler solver for sampling.

Evaluation Metrics Following existing work (Zhang et al. 2024; Yin et al. 2025), we use FVD (Ge et al. 2024) and standard VBench-I2V benchmark (Huang et al. 2024) to evaluate our model. FVD measures the Fréchet distance between

Method	Step	FVD↓	VBench			
			Temporal Quality	Frame Quality	I2V Subject	I2V Background
Teacher	25	98.06	86.83	61.92	98.72	98.12
	8	150.82	84.95	58.76	98.68	98.10
	4	321.95	81.72	51.96	98.59	98.04
LCM	4	192.93	83.35	57.51	98.61	97.07
LCM	2	591.51	79.6	45.9	98.45	97.94
PCM	4	198.56	83.23	57.29	98.61	97.05
DMD2	4	167.42	83.38	57.35	98.63	98.08
OSV [†]	4	146.59	83.97	59.79	98.75	98.14
OSV [†]	2	471.38	80.41	49.09	98.55	98.01
Ours	4	118.89	85.27	61.76	98.82	98.20
Ours	2	199.61	83.17	59.16	98.77	98.14

Table 1: Image-to-Video quantitative results on OpenVid-1M test set. FVD is computed using 64 frames. † donates our implementations.

feature distributions of real and generated videos. VBench is a comprehensive benchmark that is widely used to evaluate the video quality.

Comparisons on Image-to-Video Generation

Quantitative Comparisons. We conduct comprehensive evaluations by comparing our model against state-of-the-art distillation methods, including trajectory-preserving methods LCM (Luo et al. 2023) and PCM (Wang et al. 2024), and distribution-matching methods OSV (Mao et al. 2025) and DMD2 (Yin et al. 2024a). We implement these methods on Wan 2.1 following the official codebases of LCM, PCM, and DMD2. And we try our best to implement OSV, which we denote as OSV[†]. During evaluation, we utilize the first frame from the Openvid-1M test set as conditional input to generate 1000 videos at 832×480 resolution with 24 FPS, then compute FVD and VBench metrics on these generated videos.

As shown in Table 1, our method significantly outperforms existing state-of-the-art distillation methods across all these quantitative metrics. Specifically, our approach achieves state-of-the-art FVD scores of 118.89 with 4-step inference and 199.61 with 2-step inference, indicating that our generated videos have the closest distribution to the ground truth. In VBench evaluations, our method with 2-step inference outperforms the 4-step performance of LCM, PCM, and DMD2 in terms of frame quality. Additionally, our method with 4-step inference approaches the teacher model’s 25-step performance in frame quality, while even surpassing it in image-conditioning scores, indicating the capability of our method to preserve high-quality video synthesis under few-step inference.

Qualitative Comparisons Figure 5 presents visual comparison results of the teacher model with 8 sampling steps, LCM (Luo et al. 2023), PCM (Wang et al. 2024), OSV (Mao et al. 2025), DMD2 (Yin et al. 2024a) and our method with




Figure 4: Results from our user study regarding frame quality, temporal consistency, and image alignment of different models. The teacher method uses 8 sampling steps, while other methods use 4 steps.

4 inference steps. Figure 5 demonstrates that our method generates videos with superior visual clarity, richer details, and enhanced realism compared to other approaches, while exhibiting fewer artifacts. Additional qualitative results are provided in the appendix.

To further evaluate the performance of our method, we conduct a comprehensive user study comparing it against existing distillation approaches and the teacher model. We select first frames from 50 videos across different categories in the OpenVid-1M test set as conditional inputs. For each image, we generate 5 videos using same random seeds, with distillation methods using 4 sampling steps and the teacher model using 8 steps. Participants evaluate and select the best video based on frame quality, temporal consistency, and image alignment. As shown in Figure 4, our approach significantly outperforms baseline distillation methods and the teacher model with 8 sampling steps.

Ablation Studies

Effect of Continuous-time Consistency Distillation (CCD) To validate continuous-time consistency distillation effectiveness, we evaluate three training strategies: direct distribution alignment (DA), discrete-time consistency distillation (DCD) followed by DA, and continuous-time consistency distillation (CCD) followed by DA. Table 2a demonstrates that consistency distillation combined with distribution alignment achieves superior performance. DA alone produces excessive trajectory deviation from the original model. Furthermore, CCD outperforms DCD through more accurate ODE trajectory preservation.

Effect of Distribution Alignment (DA) To validate the effectiveness of distribution alignment, we conduct ablation studies using only continuous-time consistency loss for distillation. As shown in Table 2b, DA yields superior performance by directly approximating the real data distribution, which enhances visual fidelity under few-step sampling.

Effect of Trajectory Alignment (TA) Comparing distilled models with and without trajectory alignment, Table 2c demonstrates that model refined through TA generated higher-quality videos. This improvement stems from using DPO to align the model with high-step inference trajectories, thereby enhancing low-step inference quality.

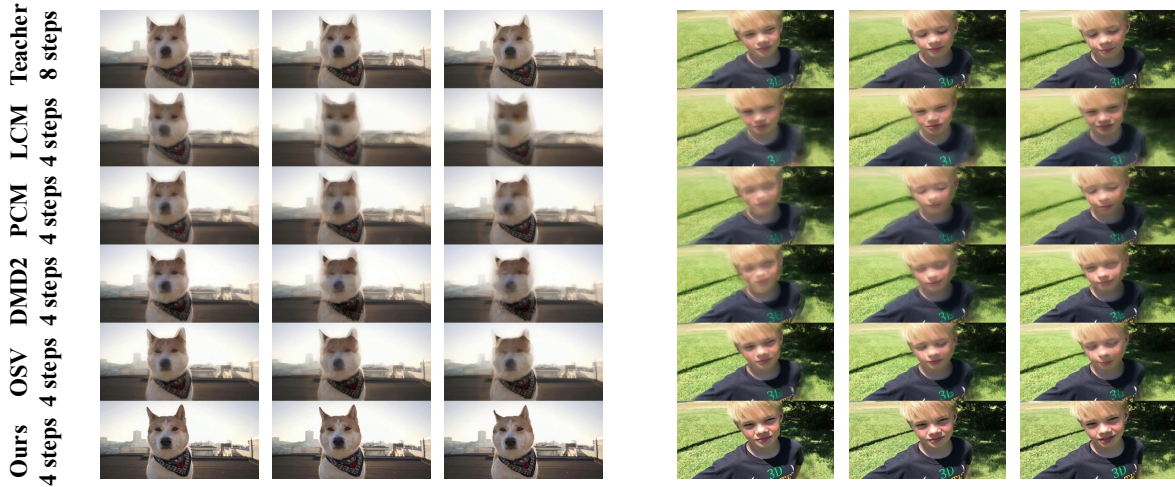


Figure 5: Qualitative comparison results. Our results significantly outperform other methods in both frame quality and temporal consistency under 4 sampling steps. All visual results are generated using images from the OpenVid-1M (Nan et al. 2025) test set.




Figure 6: Effect of synthetic preference dataset and reflow loss in trajectory alignment.

Time Samplers Prior work (Esser et al. 2024) demonstrates that the distribution of timestep affects generation quality. We train models using CCD with DA under different timestep distributions. As shown in Table 2d, the logit-normal timestep sampler achieves superior results.

Effect of Synthetic Preference Dataset in Trajectory Alignment We compared two preference dataset strategies: synthetic preference datasets versus using real data as preferred samples with low-step generated data as less preferred samples. As shown in Figure 6, synthetic preference datasets enable rapid improvement in few-step inference quality (red line), while real data as preferred samples leads to slow convergence (blue line). This occurs because synthetic preferred data implicitly contains trajectory information, allowing the model to quickly enhance low-step generation by aligning with high-step inference trajectories.

Effect of Reflow Loss in Trajectory Alignment We investigated the impact of reflow loss by comparing models trained using \mathcal{L}_{dpo} alone versus those trained with the combined objective $\mathcal{L}_{\text{DPO}} + \lambda_{\text{RF}} \cdot \mathcal{L}_{\text{RF}}$. As shown in Figure 6 (red line vs. green line), incorporating reflow loss enhances peak performance and stabilizes training.

Train setting	FVD↓	Train setting	FVD↓
DA	163.54	CCD	172.54
DCD + DA	151.23	CCD + DA	136.14
CCD + DA	136.14		

(a)

(b)

TA	FVD↓
✗	136.14
✓	118.89

(c)

t sampler	FVD↓
uniform(0, 1)	140.36
lognorm(-0.8, 1.0)	138.5
lognorm(-0.6, 1.4)	136.14

(d)

Table 2: **Ablation studies.** We set the sampling steps to 4. a) Effect of CCD. b) Effect of DA. c) Effect of TA. d) The time schedule.

Conclusion

In this paper, we propose *SwiftVideo*, a unified distillation framework that overcomes the limitations of prior approaches. Through continuous-time consistency distillation, we maintain precise ODE trajectory preservation while eliminating discretization errors. The incorporation of direct distribution alignment on real data enhances the visual fidelity of generated videos and overcomes teacher model performance constraints. Additionally, trajectory alignment across different inference steps serves as an effective post-training strategy to improve generation quality under few-step settings. Extensive experiments on OpenVid-1M demonstrate that *SwiftVideo* significantly outperforms state-of-the-art distillation methods on both FVD and VBench metrics. Future work will focus on extending this framework to larger-scale video models with even fewer steps.

Acknowledgments

The article was supported by the project from the Ministry of Science and Higher Education of the Republic of Kazakhstan, No. BR24992975 ‘Development of a digital twin of a food processing enterprise using artificial intelligence and IIoT technologies’ (2024–2026).

References

- Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendelevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; Jampani, V.; and Rombach, R. 2023. Stable Video Diffusion: Scaling Latent Video Diffusion Models to Large Datasets. arXiv:2311.15127.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Ge, S.; Mahapatra, A.; Parmar, G.; Zhu, J.-Y.; and Huang, J.-B. 2024. On the content bias in fréchet video distance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7277–7288.
- Guo, Y.; Yang, C.; Rao, A.; Liang, Z.; Wang, Y.; Qiao, Y.; Agrawala, M.; Lin, D.; and Dai, B. 2024. AnimateDiff: Animate Your Personalized Text-to-Image Diffusion Models without Specific Tuning. arXiv:2307.04725.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Huang, Z.; He, Y.; Yu, J.; Zhang, F.; Si, C.; Jiang, Y.; Zhang, Y.; Wu, T.; Jin, Q.; Chanpaisit, N.; et al. 2024. Vbench: Comprehensive benchmark suite for video generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21807–21818.
- Kong, W.; Tian, Q.; Zhang, Z.; Min, R.; Dai, Z.; Zhou, J.; Xiong, J.; Li, X.; Wu, B.; Zhang, J.; Wu, K.; Lin, Q.; Yuan, J.; Long, Y.; Wang, A.; Wang, A.; Li, C.; Huang, D.; Yang, F.; Tan, H.; Wang, H.; Song, J.; Bai, J.; Wu, J.; Xue, J.; Wang, J.; Wang, K.; Liu, M.; Li, P.; Li, S.; Wang, W.; Yu, W.; Deng, X.; Li, Y.; Chen, Y.; Cui, Y.; Peng, Y.; Yu, Z.; He, Z.; Xu, Z.; Zhou, Z.; Xu, Z.; Tao, Y.; Lu, Q.; Liu, S.; Zhou, D.; Wang, H.; Yang, Y.; Wang, D.; Liu, Y.; Jiang, J.; and Zhong, C. 2025. HunyuanVideo: A Systematic Framework For Large Video Generative Models. arXiv:2412.03603.
- Li, J.; Feng, W.; Fu, T.-J.; Wang, X.; Basu, S.; Chen, W.; and Wang, W. Y. 2024. T2v-turbo: Breaking the quality bottleneck of video consistency model with mixed reward feedback. *Advances in neural information processing systems*, 37: 75692–75726.
- Lin, S.; Wang, A.; and Yang, X. 2024. Sdxl-lightning: Progressive adversarial diffusion distillation. *arXiv preprint arXiv:2402.13929*.
- Lin, S.; Xia, X.; Ren, Y.; Yang, C.; Xiao, X.; and Jiang, L. 2025. Diffusion Adversarial Post-Training for One-Step Video Generation. arXiv:2501.08316.
- Lipman, Y.; Chen, R. T. Q.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow Matching for Generative Modeling. arXiv:2210.02747.
- Liu, J.; Liu, G.; Liang, J.; Yuan, Z.; Liu, X.; Zheng, M.; Wu, X.; Wang, Q.; Qin, W.; Xia, M.; Wang, X.; Liu, X.; Yang, F.; Wan, P.; Zhang, D.; Gai, K.; Yang, Y.; and Ouyang, W. 2025. Improving Video Generation with Human Feedback. arXiv:2501.13918.
- Liu, X.; Gong, C.; and Liu, Q. 2022. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow. arXiv:2209.03003.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. arXiv:1711.05101.
- Lu, C.; and Song, Y. 2025. Simplifying, Stabilizing and Scaling Continuous-Time Consistency Models. arXiv:2410.11081.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in neural information processing systems*, 35: 5775–5787.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2025. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *Machine Intelligence Research*, 1–22.
- Luo, S.; Tan, Y.; Huang, L.; Li, J.; and Zhao, H. 2023. Latent Consistency Models: Synthesizing High-Resolution Images with Few-Step Inference. arXiv:2310.04378.
- Mao, X.; Jiang, Z.; Wang, F.-Y.; Zhang, J.; Chen, H.; Chi, M.; Wang, Y.; and Luo, W. 2025. Osv: One step is enough for high-quality image to video generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 12585–12594.
- Nan, K.; Xie, R.; Zhou, P.; Fan, T.; Yang, Z.; Chen, Z.; Li, X.; Yang, J.; and Tai, Y. 2025. OpenVid-1M: A Large-Scale High-Quality Dataset for Text-to-video Generation. arXiv:2407.02371.
- OpenAI. 2024. Sora. <https://openai.com/sora>. "Accessed: 2024-12-09".
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; Assran, M.; Ballas, N.; Galuba, W.; Howes, R.; Huang, P.-Y.; Li, S.-W.; Misra, I.; Rabbat, M.; Sharma, V.; Synnaeve, G.; Xu, H.; Jegou, H.; Mairal, J.; Labatut, P.; Joulin, A.; and Bojanowski, P. 2024. DINOv2: Learning Robust Visual Features without Supervision. arXiv:2304.07193.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36: 53728–53741.
- Ren, Y.; Xia, X.; Lu, Y.; Zhang, J.; Wu, J.; Xie, P.; Wang, X.; and Xiao, X. 2024. Hyper-sd: Trajectory segmented consistency model for efficient image synthesis. *Advances in Neural Information Processing Systems*, 37: 117340–117362.
- Salimans, T.; and Ho, J. 2022. Progressive Distillation for Fast Sampling of Diffusion Models. arXiv:2202.00512.

Sauer, A.; Boesel, F.; Dockhorn, T.; Blattmann, A.; Esser, P.; and Rombach, R. 2024a. Fast high-resolution image synthesis with latent adversarial diffusion distillation. In *SIG-GRAPH Asia 2024 Conference Papers*, 1–11.

Sauer, A.; Karras, T.; Laine, S.; Geiger, A.; and Aila, T. 2023. Stylegan-t: Unlocking the power of gans for fast large-scale text-to-image synthesis. In *International conference on machine learning*, 30105–30118. PMLR.

Sauer, A.; Lorenz, D.; Blattmann, A.; and Rombach, R. 2024b. Adversarial diffusion distillation. In *European Conference on Computer Vision*, 87–103. Springer.

Wan, T.; Wang, A.; Ai, B.; Wen, B.; Mao, C.; Xie, C.-W.; Chen, D.; Yu, F.; Zhao, H.; Yang, J.; Zeng, J.; Wang, J.; Zhang, J.; Zhou, J.; Wang, J.; Chen, J.; Zhu, K.; Zhao, K.; Yan, K.; Huang, L.; Feng, M.; Zhang, N.; Li, P.; Wu, P.; Chu, R.; Feng, R.; Zhang, S.; Sun, S.; Fang, T.; Wang, T.; Gui, T.; Weng, T.; Shen, T.; Lin, W.; Wang, W.; Wang, W.; Zhou, W.; Wang, W.; Shen, W.; Yu, W.; Shi, X.; Huang, X.; Xu, X.; Kou, Y.; Lv, Y.; Li, Y.; Liu, Y.; Wang, Y.; Zhang, Y.; Huang, Y.; Li, Y.; Wu, Y.; Liu, Y.; Pan, Y.; Zheng, Y.; Hong, Y.; Shi, Y.; Feng, Y.; Jiang, Z.; Han, Z.; Wu, Z.-F.; and Liu, Z. 2025. Wan: Open and Advanced Large-Scale Video Generative Models. arXiv:2503.20314.

Wang, F.-Y.; Huang, Z.; Bergman, A.; Shen, D.; Gao, P.; Lingelbach, M.; Sun, K.; Bian, W.; Song, G.; Liu, Y.; et al. 2024. Phased consistency models. *Advances in neural information processing systems*, 37: 83951–84009.

Wang, X.; Zhang, S.; Zhang, H.; Liu, Y.; Zhang, Y.; Gao, C.; and Sang, N. 2023. VideoLCM: Video Latent Consistency Model. arXiv:2312.09109.

Yang, Z.; Teng, J.; Zheng, W.; Ding, M.; Huang, S.; Xu, J.; Yang, Y.; Hong, W.; Zhang, X.; Feng, G.; Yin, D.; Zhang, Y.; Wang, W.; Cheng, Y.; Xu, B.; Gu, X.; Dong, Y.; and Tang, J. 2025. CogVideoX: Text-to-Video Diffusion Models with An Expert Transformer. arXiv:2408.06072.

Yin, T.; Gharbi, M.; Park, T.; Zhang, R.; Shechtman, E.; Durand, F.; and Freeman, B. 2024a. Improved distribution matching distillation for fast image synthesis. *Advances in neural information processing systems*, 37: 47455–47487.

Yin, T.; Gharbi, M.; Zhang, R.; Shechtman, E.; Durand, F.; Freeman, W. T.; and Park, T. 2024b. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6613–6623.

Yin, T.; Zhang, Q.; Zhang, R.; Freeman, W. T.; Durand, F.; Shechtman, E.; and Huang, X. 2025. From slow bidirectional to fast autoregressive video diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 22963–22974.

Zhang, Z.; Li, Y.; Wu, Y.; Kag, A.; Skorokhodov, I.; Mena-pace, W.; Siarohin, A.; Cao, J.; Metaxas, D.; Tulyakov, S.; et al. 2024. Sf-v: Single forward video generation model. *Advances in Neural Information Processing Systems*, 37: 103599–103618.