DATA-FORCING DISTILLATION: RESTORING DIVERSITY AND FIDELITY IN FEW-STEP VIDEO GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Recent progress has shown promise in distilling multi-step video diffusion models into efficient few-step students. Among them, Distribution Matching Distillation (DMD) and its successor DMD2 achieve strong generation quality and fast convergence. However, due to the nature of the reverse Kullback–Leibler (KL) objective, these methods can exhibit two persistent failure modes: reduced sample diversity and visibly over-saturated outputs that deviate from real-video appearance. In this work, we propose *Data-Forcing Distillation* (DFD), a simple post-training framework that improves DMD-based students *with only a single line of code change*. At its core is the *teacher score discrepancy*, which guides the student toward the reference-data distribution by pulling probability mass toward underrepresented modes while suppressing problematic modes that are inconsistent with the reference data. We provide an in-depth theoretical analysis of our framework and evaluate it on text-to-video, image-to-video, and autoregressive video generation. With only *100–300 steps* of finetuning, DFD improves both sample diversity and visual quality in text-to-video and autoregressive generation, while improving visual quality and temporal coherence in image-to-video generation. These results show that DFD can mitigate key failure modes of few-step video generators across different generation settings. We refer readers to the supplementary material for full video comparisons.

1 INTRODUCTION

Recent progress in large-scale diffusion and flow-based models has substantially advanced video generation Song et al. (2021); Lipman et al. (2023); Rombach et al. (2022); Wan et al. (2025); NVIDIA et al. (2026); Kong et al. (2025); Yang et al. (2025); Brooks et al. (2024). Modern video generators can synthesize photorealistic, high-fidelity, and diverse videos from a text prompt or a single image. Yet, this capability comes at a heavy computational cost: high-quality generation typically requires multiple denoising steps, resulting in high inference latency and limited deployment in interactive, real-time, or large-scale production settings. Accelerating diffusion sampling while preserving the quality and diversity of the original multi-step model has thus become a central research challenge.

Two families of distillation methods have emerged to address this challenge. *Trajectory-based* methods train the few-step student to regress the sampling trajectory of a multi-step teacher Salimans & Ho (2022); Meng et al. (2023); Song et al. (2023); Song & Dhariwal (2024); Lu & Song (2025); Geng et al. (2026); Sabour et al.. While effective at reducing sampling steps, these methods require the student to approximate a complex multi-step trajectory with only a few function evaluations, which becomes challenging for high-dimensional video generation and can result in noticeable quality degradation. *Distribution-based* methods Yin et al. (2024a;b); Luo et al. (2023b); Wang et al. (2026) relax this trajectory constraint and train the student so that its output distribution matches the teacher's distribution. DMD Yin et al. (2024a) and its successor DMD2 Yin et al. (2024b) are representative examples, achieving strong few-step quality and fast convergence; we therefore build on this family in our work. However, when applied to large-scale video generation, distribution-based methods still exhibit two critical failure modes: a substantial drop in sample diversity compared to the teacher model, and visibly over-saturated outputs that deviate from real-video appearance.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

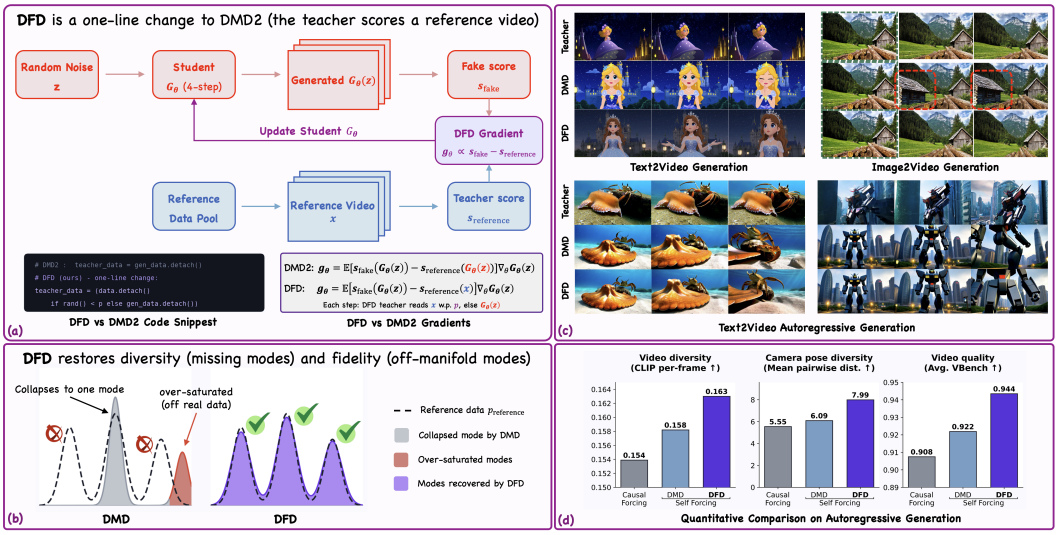

Figure 1: **Overview of DFD.** (a) DFD is a one-line change to DMD2, where the teacher scores a reference video x (w.p. p) instead of the student's generation $G_\theta(z)$, giving the update $g_\theta \propto s_{\text{fake}} - s_{\text{reference}}$. (b) This counters reverse-KL mode-seeking, recovering collapsed (diversity) and over-saturated off-manifold (fidelity) modes. (c) Qualitative results on text-to-video, image-to-video, and autoregressive text-to-video generation: DFD avoids over-saturation, follows the input frame, and keeps objects coherent over long horizons. (d) On autoregressive generation, DFD improves video diversity, camera-pose diversity, and average VBench quality over Causal Forcing and Self Forcing (DMD).

One of the main cause of both failure modes is the reverse Kullback-Leibler (KL) objective used in DMD, which is well known to be *mode-seeking* Poole et al. (2022); Wang et al. (2023): the student seeks the highest-density modes of p_{teacher}, as the reverse KL takes its expectation over student-generated samples (Sec. 2) and primarily optimizes only the modes that the student covers and cannot penalize missing modes in the teacher distribution.

To fix the problem, our intuition is simple: directly incorporating reference video samples into the training objective during distillation, so that the student can be explicitly pulled toward the diverse, high-quality reference data distribution rather than supervising itself using its own generated data.

Building on this motivation, we propose *Data-Forcing Distillation* (DFD), a video distillation framework that restores both diversity and fidelity with a single line of code change on top of DMD. The core of DFD is the *teacher score discrepancy* that computes the score difference between a reference video and the student's generated video, and we directly incorporate it in the distillation objective. Intuitively, this score discrepancy points the student to the reference-data distribution, pulling it toward missing modes (mitigating mode collapse) and away from problematic modes absent in reference data (avoiding over-saturation).

In practice, DFD is implemented with a single-line code change: feeding a reference video sample to the teacher in place of the student's own generation when computing the distribution matching objective. While the idea of leveraging reference data was already explored in DMD2 Yin et al. (2024b), it was implemented through an auxiliary GAN Goodfellow et al. (2014), which destabilizes training and only *implicitly* transfers reference-data information to the generator through a discriminator. In contrast, DFD *explicitly* provides reference-data guidance to the distribution matching objective, preserving the efficiency and simplicity of DMD while mitigating both failure modes.

We conduct extensive experiments on three different video generation tasks with different pretrained video models to demonstrate the effectiveness of DFD. On text-to-video generation and autoregressive video generation, we distill the pretrained Wan2.1-1.3B Wan et al. (2025) model into a 4-step bidirectional/causal student model. On image-to-video generation, we apply the same recipe and distill the pretrained Cosmos-Predict2.5-2B NVIDIA et al. (2026) model. With only 100-300 steps of post-training on top of DMD2, our DFD consistently and significantly improves visual quality,

restores diversity across scene composition and camera motion quality of DMD2, and substantially alleviates over-saturation artifacts. We will release our code and model checkpoints to support future research on efficient, high-fidelity, and diverse video generation.

2 Background: Distribution Matching Distillation

Diffusion distillation broadly splits into two families: *trajectory-based* methods, which train the student to follow the teacher's sampling trajectory Song et al. (2023); Lu & Song (2025); Zheng et al. (2026), and *distribution-matching* methods, which instead align the student's output distribution with the teacher's Yin et al. (2024a;b); Xu et al. (2025); Wang et al. (2026); Wu et al. (2026); Nie et al. (2026). We build on the latter, whose representative instances are DMD Yin et al. (2024a) and DMD2 Yin et al. (2024b); a comprehensive review of both families is provided in Appendix A. We first review Distribution Matching Distillation (DMD), the reverse-KL objective it optimizes, and the mode-collapse issue that motivates our approach.

Distribution Matching Distillation (DMD) Yin et al. (2024a;b). To distill a pretrained multi-step teacher into a few-step student G_θ, DMD matches the student's output distribution p_{fake} to the teacher's data distribution p_{teacher} via a KL objective. Since neither density is tractable, DMD reformulates the gradient in terms of *score functions*, which can be approximated by diffusion denoisers. Concretely, for a noise-perturbed sample $\boldsymbol{x}_t = \alpha_t G_\theta(\boldsymbol{z}, c) + \sigma_t \boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$, the student gradient takes the form

$$\nabla_\theta \mathcal{L}_{\text{DMD}} \;=\; \mathbb{E}_{t,\boldsymbol{z},\boldsymbol{\epsilon}}\big[\, w(t)\,\big(\nabla_{\boldsymbol{x}} \log p_{\text{fake}}(\boldsymbol{x}_t) - \nabla_{\boldsymbol{x}} \log p_{\text{teacher}}(\boldsymbol{x}_t)\big)\,\nabla_\theta G_\theta(\boldsymbol{z}, c)\big]\,, \tag{1}$$

where $\nabla_{\boldsymbol{x}} \log p_{\text{teacher}}$ is learned by the frozen teacher denoiser, $\nabla_{\boldsymbol{x}} \log p_{\text{fake}}$ is learned by an auxiliary denoiser trained *online* on student samples via denoising score matching, and $w(t)$ is a noise-level weighting. Training alternates between updating θ with Eq. (1) and updating the auxiliary denoiser to track the moving p_{fake}.

Reverse Kullback-Leibler (KL) Loss. DMD instantiates Eq. (1) as the gradient of the *reverse* KL,

$$D_{\text{KL}}\big(p_{\text{fake}} \,\|\, p_{\text{teacher}}\big) \;=\; \mathbb{E}_{\substack{\boldsymbol{x}=G_\theta(\boldsymbol{z},c)\\ \boldsymbol{z}\sim\mathcal{N}(\mathbf{0},\boldsymbol{I})}}\left[\log \frac{p_{\text{fake}}(\boldsymbol{x})}{p_{\text{teacher}}(\boldsymbol{x})}\right], \tag{2}$$

along the diffusion path whose differentiation yields

$$\nabla_\theta D_{\text{KL}}(p_{\text{fake}} \,\|\, p_{\text{teacher}}) \;=\; \mathbb{E}_{\boldsymbol{z}}\big[\big(\nabla_{\boldsymbol{x}} \log p_{\text{fake}}(\boldsymbol{x}) - \nabla_{\boldsymbol{x}} \log p_{\text{teacher}}(\boldsymbol{x})\big)\,\nabla_\theta G_\theta(\boldsymbol{z}, c)\big]\,, \quad \boldsymbol{x} = G_\theta(\boldsymbol{z}, c). \tag{3}$$

This formulation gives strong few-step quality and fast convergence, but the reverse KL is well known to be *mode-seeking*: the student collapses onto the highest-density modes of p_{teacher}, producing low-diversity, often over-saturated samples He et al. (2024); Lu et al. (2025). The core reason for this mode-seeking behavior is that reverse KL divergence is evaluated as an expectation over the student's generated distribution. Consequently, if the student drops a mode, the loss function incurs no penalty in that region, yielding zero gradient signal to recover the missing mode. The core reason for mode-seeking is that reverse KL samples from the data generated by the student, and when there is a missing mode in the student, the KL divergence can not pull it back.

Regularization with Reference Data To overcome the stated mode collapse, over-saturated problem introduced by reverse KL loss, DMD2 Yin et al. (2024b) partially compensates with an auxiliary GAN Goodfellow et al. (2014) discriminator, but adversarial training is unstable and only *implicitly* injects the reference-data signal through the discriminator rather than through the distribution-matching gradient itself. Adversarial Diffusion Distillation(ADD) Sauer et al. (2024) likewise injects reference-data supervision through an auxiliary adversarial discriminator, while its frozen diffusion teacher is still evaluated on noised *student-generated* samples, as in standard score distillation. DFD uses no auxiliary discriminator; it changes the teacher-score *input* directly, evaluating the teacher on a noised reference sample during DFD steps and on a noised generated sample during DMD steps. The main distinction is thus *where* and *how* the reference data enters the distillation gradient.

3 Data-Forcing Distillation

Our core intuition is to incorporate reference data directly into the reverse KL gradient through a differentiable regularizer, where the student's distribution can be explicitly pulled toward the **diverse, high-quality reference data distribution** rather than supervising itself in the teacher. In this section, we first formally formulate our proposed data-forcing distillation, then give a practical implementation.

3.1 THE DATA-FORCING DISTILLATION (DFD) FRAMEWORK

Our DFD injects a reference-data regularization term into the original DMD gradient(we use p_{refer} to indicate the distribution of reference data).

$$\begin{aligned} g_{\text{DFD}}(\theta) \;=\; & \underbrace{\mathbb{E}_{\boldsymbol{z}\sim\mathcal{N}(\mathbf{0},\boldsymbol{I})}\Big[\big(\nabla_{\boldsymbol{x}}\log p_{\text{fake}}(G_\theta(\boldsymbol{z},c)) - \nabla_{\boldsymbol{x}}\log p_{\text{refer}}(G_\theta(\boldsymbol{z},c))\big)\,\nabla_\theta G_\theta(\boldsymbol{z},c)\Big]}_{\text{native DMD gradient } g_{\text{DMD}}(\theta)} \\ & - \underbrace{\mathbb{E}_{\boldsymbol{x}\sim p_{\text{refer}}(\cdot|c),\,\boldsymbol{z}\sim\mathcal{N}(\mathbf{0},\boldsymbol{I})}\Big[\Delta_{p_{\text{refer}},\,\boldsymbol{x},\,G_\theta(\boldsymbol{z})}\,\nabla_\theta G_\theta(\boldsymbol{z},c)\Big]}_{\text{reference-data regularizer}}. \end{aligned} \tag{4}$$

$$\Delta_{p_{\text{refer}},\boldsymbol{x},G_\theta(\boldsymbol{z})} \;=\; \nabla_{\boldsymbol{x}}\log p_{\text{refer}}(\boldsymbol{x}) - \nabla_{\boldsymbol{x}}\log p_{\text{refer}}(G_\theta(\boldsymbol{z},c)), \quad \boldsymbol{x}\sim p_{\text{refer}}(\cdot\mid c). \tag{5}$$

We refer to $\Delta_{p_{\text{refer}},\boldsymbol{x},G_\theta(\boldsymbol{z})}$ as the **teacher score discrepancy**: it measures the gap, in the teacher's score field, between a reference sample and the student's generation. When the student matches the reference distribution ($p_{\text{fake}} = p_{\text{refer}}$), this discrepancy is minimized to zero in expectation, $\mathbb{E}_{\boldsymbol{x},\boldsymbol{z}}[\Delta_{p_{\text{refer}},\boldsymbol{x},G_\theta(\boldsymbol{z})}] = 0$ (proof in the Appendix D). Empirically, drawing $\boldsymbol{x}$ from diverse, high-quality reference data and minimizing $\Delta_{p_{\text{refer}},\boldsymbol{x},G_\theta(\boldsymbol{z})}$ pulls the student toward modes it has missed in the reference data (mitigating mode collapse) and away from problematic modes such as over-saturated outputs that are not in the reference data. More importantly, properties of reference data, such as temporal coherence, physical plausibility, photorealism, can be explicitly distilled into the student through this term.

Discussion: No Need for a GAN. DMD2 Yin et al. (2024b) provides *implicit* reference-data supervision through an auxiliary GAN loss; in contrast, DFD provides *explicit* reference-data supervision through the score-discrepancy term (Eq. 5). As a result, we remove the GAN loss entirely from our distillation pipeline, leading to substantially improved in video dynamics while maintaining comparable performance across the other evaluation metrics. We validate this in the ablation study in Appendix C.

3.2 THE STABILITY OF DFD

The teacher-score discrepancy $\Delta_{p_{\text{refer}},\boldsymbol{x},G_\theta(\boldsymbol{z},c)}$ is well behaved in expectation. Our goal is to ensure that the DFD objective remains a controlled perturbation of the original DMD objective. While we evaluate the score along the full diffusion path ($\boldsymbol{x}_t = \alpha_t\boldsymbol{x} + \sigma_t\boldsymbol{\epsilon}$, both for the reference sample and the student's generation) rather than only at its endpoint, we also introduce the validity conditon:

$$\mathbb{E}_{\boldsymbol{x},\boldsymbol{z}}\Big[\|\boldsymbol{x} - G_\theta(\boldsymbol{z},c)\|_2^2 \;\mid c\Big] \;\le\; \delta(c)^2, \tag{6}$$

Further assuming the second-moment generator-Jacobian bound $\mathbb{E}_{\boldsymbol{z}}[\|\nabla_\theta G_\theta(\boldsymbol{z},c)\|_{\text{op}}^2 \mid c] \le B^2$, Cauchy–Schwarz yields

$$\begin{aligned} \|g_{\text{DFD}}(\theta) - g_{\text{DMD}}(\theta)\|_2^2 &= \big\|\mathbb{E}_{\boldsymbol{x},\boldsymbol{z}}\big[\Delta_{p_{\text{refer}},\boldsymbol{x},G_\theta(\boldsymbol{z},c)}\nabla_\theta G_\theta(\boldsymbol{z},c)\mid c\big]\big\|_2^2 \\ &\le B^2\alpha_t^2 L(c,t)^2\,\mathbb{E}_{\boldsymbol{x},\boldsymbol{z}}\Big[\|\boldsymbol{x} - G_\theta(\boldsymbol{z},c)\|_2^2 \mid c\Big] \\ &\le B^2\alpha_t^2 L(c,t)^2\delta(c)^2. \end{aligned} \tag{7}$$

where $L(c,t)$ is the Lipschitz constant of the teacher's score network. Thus, when the validity condition holds and the score sensitivity and generator Jacobian are controlled, DFD remains a bounded perturbation of the DMD update. The complete derivation is provided in Appendix D.

Discussion: Alignment Between the Reference Sample and the Generated Sample The alignment required by DFD is *distributional, not pairwise*. According to Eq. 6, we do not assume that an individual reference sample $\boldsymbol{x}$ and an individual generated sample $G_\theta(\boldsymbol{z},c)$ are semantically or geometrically matched merely because they share the same condition c. Accordingly, the teacher score evaluated at $\boldsymbol{x}$ should not be interpreted as a deterministic, instance-level transport direction from $\boldsymbol{x}$ to $G_\theta(\boldsymbol{z},c)$. Instead, $\boldsymbol{x}\sim p_{\text{refer}}(\cdot\mid c)$ and $G_\theta(\boldsymbol{z},c)\sim p_{\text{fake}}(\cdot\mid c)$ provide stochastic samples from the two conditional distributions, and our analysis concerns the resulting update *in expectation*. The additional discussion is in Appendix D.2

3.3 Practical Implementation

Eq. (4) admits a clean simplification. Because the regularizer $\Delta_{p_{\text{refer}},\boldsymbol{x},G_\theta(\boldsymbol{z})}$ is evaluated at the same noise level and condition c as the DMD gradient, it exactly cancels the score $\nabla_{\boldsymbol{x}} \log p_{\text{refer}}(G_\theta(\boldsymbol{z}, c))$. The DFD gradient therefore reduces to

$$g_{\text{DFD}}(\theta) \;=\; \mathbb{E}_{\substack{\boldsymbol{x}\sim p_{\text{refer}}(\cdot|c)\\ \boldsymbol{z}\sim\mathcal{N}(\mathbf{0},\boldsymbol{I})}}\left[\left(\nabla_{\boldsymbol{x}} \log p_{\text{fake}}(G_\theta(\boldsymbol{z}, c)) - \nabla_{\boldsymbol{x}} \log p_{\text{refer}}(\boldsymbol{x})\right) \nabla_\theta G_\theta(\boldsymbol{z}, c)\right]. \tag{8}$$

The only difference compared to the DMD gradient (Eq. 3) is the second score: instead of evaluating $\nabla_{\boldsymbol{x}} \log p_{\text{refer}}$ at the student's own output $G_\theta(\boldsymbol{z}, c)$, we evaluate it at a reference sample $\boldsymbol{x}$ drawn under the same condition. Fig. 2 illustrates this difference. In practice, we do not enforce Eq. (6) as a hard constraint. Instead, we satisfy it implicitly by applying DFD as a post-training stage on top of a DMD2-pretrained model, whose generations are already close to the reference videos and naturally keep $\mathbb{E}_{\boldsymbol{x},\boldsymbol{z}}[\|\boldsymbol{x} - G_\theta(\boldsymbol{z}, c)\|_2^2]$ small. The ablation in section 4.2 confirms the importance of this regime: violating the validity condition—e.g., applying DFD from scratch when the student still produces noisy videos—prevents the model from converging. Concretely, we form the practical update as a combination of the DMD gradient and the DFD gradient,

$$g(\theta) \;=\; (1-w)\, g_{\text{DMD}}(\theta) \;+\; w\, g_{\text{DFD}}(\theta), \tag{9}$$

where $w \in [0, 1]$ controls how much of the DMD signal is retained. We use $w = \frac{1}{2}$ as our default, which preserves the fast convergence of the DMD while leveraging reference-data guidance.

In practice, we implement Eq. 9 via per-step stochastic sampling, which matches in expectation when $p = w$:

$$\nabla_\theta \mathcal{L} = \begin{cases} g_{\text{DFD}}(\theta) & \text{with probability } p, \\ g_{\text{DMD}}(\theta) & \text{with probability } 1-p. \end{cases} \tag{10}$$

Additionally, following standard practice in the DMD line of work, we replace the exact score in Eq. 9 with the scores estimated by the diffusion models on perturbed samples, and take the expectation over the diffusion timesteps.

The pseudocode for DFD The pseudocode for the student update is provided in Appendix B.

3.4 Reference

Data Selection for DFD

Once reference data enters the distillation gradient, its quality and diversity become the new bottleneck on what the student can learn. We therefore curate the training videos with high quality and diversity. Specifically, we start from the public ViPE-Wild-1M dataset Huang et al. (2025a), whose videos were generated offline by the Wan 2.1 Wan et al. (2025) 14B model, and cluster its 1M videos into 1,000 clusters. We then manually retain the clusters whose videos exhibit clear semantics, temporal coherence, varied styles, and diverse dynamics. From the retained pool, we construct two training sets: a 20,000-video animation and cartoon set and a complete 30,000-video mixed-style set. We detail the exact selection process in Appendix E.

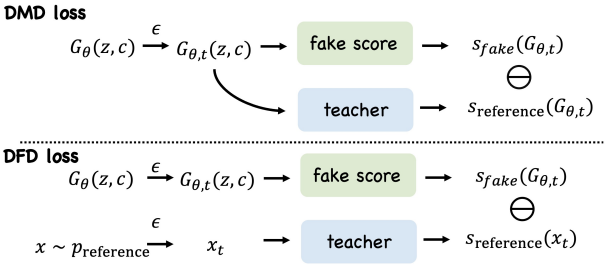

Figure 2: **The comparison between our DFD and the original DMD**. Our DFD computes the reference score directly using the videos sampled from the reference data distribution, while the original DMD computes the reference score using the generated videos from the student.

Controlling for the Data Source. All distilled baselines are trained on the same 14B-derived reference corpus as DFD; the only difference is how that corpus is used. DMD2 incorporates it *implicitly* through the GAN discriminator, whereas DFD injects it *explicitly* through the teacher-score discrepancy. Because the data source is held fixed across methods, the relative improvement is attributable to how the data is utilized rather than to how it was produced. We further isolate this effect in section 4.2, where supervised fine-tuning of the teacher on exactly the same corpus degrades quality and diversity.



Figure 3: **Qualitative results on text-to-video generation.** The left columns show models distilled on animation set, and the right column on the mix-style set. Our method produces videos that are not over-saturated, and recovers finer details such as the wing texture of the phoenix.

Table 1: Quantitative results on the text-to-video experiments using Wan2.1-1.3B. For all metrics, higher is better. The best result among the distilled methods in each column is **bolded**.

	Video Diversity		Camera Pose Diversity		Video Quality		
Method	**CLIP (Per-frame)**	**DINO (Per-frame)**	**Endpoint Distance**	**Trajectory Distance**	**Average VBench**	**VisionReward ×100**	**Instruction Following ×100**
Teacher	0.222	0.350	30.571	26.651	0.899	6.00	61.0
DMD2	0.165	0.239	9.148	4.284	0.901	5.08	56.9
DP-DMD	0.165	0.239	7.208	3.466	0.903	5.25	57.1
Ours	**0.170**	**0.252**	**18.513**	**19.256**	**0.906**	**5.85**	**59.8**

4 Experiments and Results

We evaluate DFD on three representative large-scale video-generation settings: Wan2.1-1.3B for text-to-video generation, auto-regressive video generation and Cosmos-Predict2.5-2B for image-to-video generation, We further conduct a comprehensive ablation study to demonstrate the effectiveness of our design choices.

4.1 Main Experiments

4.1.1 Text-to-Video Generation

Experimental Settings. For the text-to-video generation, we distill the Wan2.1-1.3B model on both the animation set and the mix-style set. We compare against two baselines: DMD2 Yin et al. (2024b) and DP-DMD Wu et al. (2026). For DMD2, we use the implementation and default configurations from the FastGen codebase. For DP-DMD, we follow the original paper and use a diversity anchor step of $K = 5$ with weight 0.05 Wu et al. (2026). For all the methods, we distill the pre-trained Wan model into a four-step student model. We evaluate results along three aspects. **Video quality**: Following *Causal Forcing* Zhu et al. (2026), we adopt the metrics from VBench Huang et al. (2023), VisionReward Xu et al. (2026) and its instruction-following sub-score. **Video diversity**: we adopt the diversity metric from DP-DMD Wu et al. (2026). **Camera-pose diversity**: we use the ViPE Huang et al. (2025a) to estimate the camera poses of each generated video, and compute diversity statistics from the estimated cameras. The evaluation set comprises 70 animation prompts and 70 mix-style prompts, respectively. We generate 8 videos for each prompt with random seeds 1–8, yielding 1120 videos in total for evaluation. Further details are provided in the Appendix E.

Experimental Results We provide qualitative results in Fig. 3 with quantitative results averaged on the models distilled on the animation test set and the mix-style testset in Table 1. Qualitatively, our method produces higher-quality videos: it mitigates the over-saturation artifacts with significantly better appearance and dynamics, and recovers fine-grained details. Quantitatively, DFD obtains the best average VBench, VisionReward, and instruction-following scores among the distilled models. DFD also improves both visual-diversity metrics and substantially increases camera-pose



Figure 4: **I2V results on the VBench test set.** The colored outline indicates the input image. Our method shows much better visual quality with smooth and physically plausible dynamics under complex scenarios, whereas DP-DMD and DMD2 fail to produce valid videos.

Table 2: Quantitative comparison on image-to-video generation. For all metrics, higher is better, and the best result in each column is **bolded**.

	ViPE test set†			VBench I2V set‡		
Method	**Average VBench**	**VisionReward ×100**	**Instruction Following ×100**	**Average VBench**	**VisionReward ×100**	**Instruction Following ×100**
Teacher	0.9151	9.90	52.2	0.9285	12.96	81.6
DMD2	0.9171	8.06	**49.8**	0.9245	11.92	73.9
DP-DMD	0.8789	4.97	37.3	0.8843	8.59	55.4
Ours	**0.9207**	**8.83**	49.0	**0.9303**	**12.72**	**78.9**

† Results on the held-out ViPE image-to-video test set. ‡ Results on the VBench image-to-video test set.

diversity compared with DMD2 and DP-DMD. Additional qualitative results are provided in Appendix F.1.

4.1.2 Image-to-Video Generation

Experimental Settings. For image-to-video generation, we distill the Cosmos-Predict2.5-2B model on our curated mix-style dataset. For evaluation, we use two image-to-video test sets: one from VBench, containing 348 images, and one curated from the ViPE-Wild-1M dataset, containing 78 images that are held out from training. We adopt the image-to-video metrics from VBench Huang et al. (2023) for evaluation. We do not report the diversity metrics used in the image-to-video experiments because image-to-video generation is strongly conditioned on the input frame and evaluating conditional diversity would require multiple samples per input, which is outside the scope of general image-to-video evaluation. Further details are provided in the Appendix E.

Experimental Results. We provide quantitative results on both test sets in Table 2, with qualitative results in Fig 4 and Fig 8. Quantitatively, our method outperforms the baselines on the majority of metrics. Qualitatively, DP-DMD frequently violates the first-frame constraint, occasionally producing abrupt artifacts that disrupt the whole frame. DMD2 yields visible quality improvements over DP-DMD but still struggles to maintain temporal coherence in later frames and often hallucinates structural anomalies such as distorted human figures appearing outside vehicles. In contrast, our model adheres closely to the input frame and remains stable and consistent across all generated frames. Our model also achieves better physical plausibility. In the traffic scene example shown in Fig. 4, DMD2 produces visibly implausible interactions, such as vehicles intersecting. DFD reduces such artifacts in this example and yields more coherent motion.

4.1.3 Autoregressive Video Generation

Autoregressive video diffusion generates a video chunk by chunk under causal attention, which makes long-horizon and streaming synthesis possible Chen et al. (2024). To reach interactive speed, CausVid Yin et al. (2025), Self Forcing Huang et al. (2025b), and Causal Forcing Zhu et al. (2026) distill a bidirectional teacher into a causal few-step student with distribution-matching objectives,

Table 3: **Quantitative comparison on autoregressive video generation**, evaluated on 70 prompts with eight seeds per prompt. For all metrics, higher is better. The best result among the distilled autoregressive methods in each column is **bolded**; the multi-step teacher is reported as an oracle reference.

	Video Diversity		Camera Pose Diversity		Video Quality				
Method	**CLIP (Per-frame)**	**DINO (Per-frame)**	**Endpoint Spread**	**Mean Pairwise Distance**	**Background Consistency**	**Temporal Flickering**	**Motion Smoothness**	**Dynamic Degree**	**Average VBench**
Teacher (oracle)	0.2220	0.3500	18.513	26.651	0.9590	0.9760	0.9850	0.9260	0.9615
Causal Forcing Zhu et al. (2026)	0.1539	0.2496	4.35	5.55	**0.9557**	0.9425	0.9551	0.7768	0.9075
Self Forcing (DMD) Huang et al. (2025b)	0.1582	0.2502	4.85	6.09	0.9285	0.9316	0.9542	0.8732	0.9219
Self Forcing (DFD)	**0.1630**	**0.2658**	**6.17**	**7.99**	0.9303	**0.9573**	**0.9773**	**0.9089**	**0.9435**

Table 4: **Quantitative comparison on the full 100-prompt evaluation set of Causal Forcing**, using its complete evaluation suite. For all metrics, higher is better. The best result in each column is **bolded**.

Method	**Dynamic Degree**	**VBench Quality**	**VBench Semantic**	**VisionReward Total** $\times 100$	**VR Instruction Following** $\times 100$
Causal Forcing Zhu et al. (2026)	0.91	0.6864	0.6498	9.38	53.3
Self Forcing (DMD) Huang et al. (2025b)	0.95	0.6812	0.6504	9.14	50.7
Self Forcing (DFD)	**0.98**	**0.7085**	**0.6836**	**9.62**	**55.3**

and therefore inherit the same reverse-KL mode seeking that DFD is designed to correct; see Appendix A for a fuller discussion.

Experimental Settings. For autoregressive video generation, we distill the Wan2.1-1.3B model on our curated mixed-style dataset following the *Self Forcing* framework Huang et al. (2025b). All methods are implemented within the Self Forcing codebase. We compare against two autoregressive baselines: *Self Forcing* Huang et al. (2025b), which uses the DMD objective, and *Causal Forcing* Zhu et al. (2026). We apply DFD as a post-training stage on top of the Self Forcing checkpoint. We report two evaluation protocols. The first follows the same protocol as our text-to-video experiments: 70 prompts with eight videos per prompt generated from different random seeds, scored with VBench quality metrics, the feature-space diversity metrics, and the camera-pose diversity statistics. The second follows the full evaluation suite of Causal Forcing on its complete 100-prompt evaluation set, which reports VBench quality and semantic scores together with VisionReward Xu et al. (2026) and its instruction-following sub-score. All methods are evaluated under identical settings. Additional implementation details are provided in Appendix E.

Experimental Results. Qualitative results are presented in Fig. 1 and Fig. 5. Similar to our image-to-video experiments, Self Forcing with DMD struggles to preserve temporal coherence over long horizons and frequently introduces structural artifacts and hallucinations. For example, crab-like figures gradually emerge from the stone scene and Gundam gradually breaks up. In the train example shown in Fig. 5, DMD generates physically inconsistent interactions including disconnected train carriages. DFD substantially reduces such artifacts and yields a more coherent and realistic scene. These observations echo our I2V findings and further support the core intuition behind our approach: by incorporating reference data into the distillation process, the generator not only alleviates mode-seeking behavior and over-saturation artifacts, but also encourages the model to capture key characteristics of reference videos, such as inter-frame coherence, temporal consistency, and physical realism. Quantitative results under our 70-prompt protocol are reported in Table 3. Self Forcing (DFD) outperforms both autoregressive baselines on every reported metric except background consistency. Relative to Self Forcing (DMD), DFD improves temporal flickering, motion smoothness, dynamic degree, feature-space diversity, and camera-motion diversity. The full multi-step teacher remains stronger overall, as expected, while these results show that DFD is also an effective post-training correction inside an autoregressive generation pipeline. Table 4 reports the same comparison on the full 100-prompt evaluation set of Causal Forcing using its complete evaluation suite. The trend is consistent with the 70-prompt evaluation: DFD achieves the best score on all reported metrics, confirming that its improvements carry over to the full, commonly used benchmark.

4.2 Ablation Study

Effects of the Reference Data. Because DFD draws its supervision from the curated reference corpus, a natural question is whether the gains come from the corpus itself rather than from how DFD

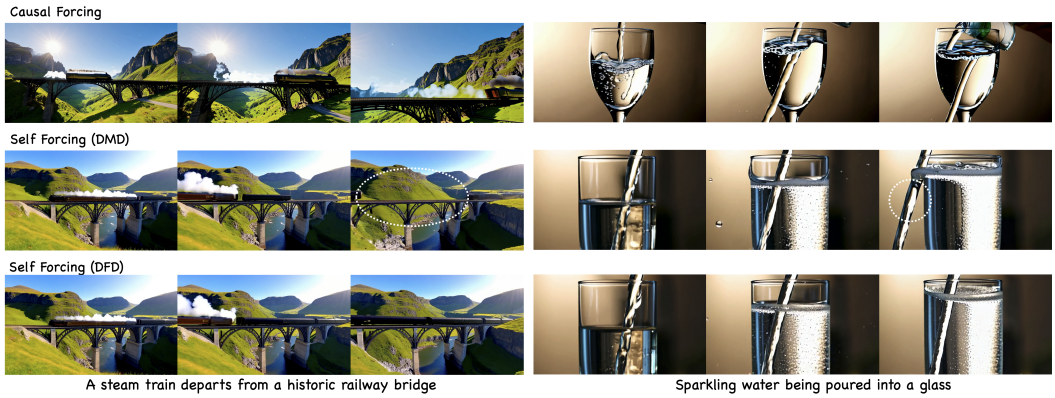

Figure 5: **Autoregressive video generation results.** Results from the generator trained with the Self Forcing pipeline using DFD loss outperform those trained with DMD loss, especially in frame consistency and physical plausibility.

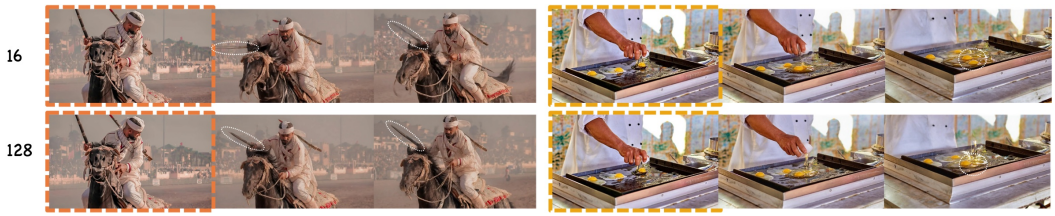

Figure 9: **Effect of scaling up the distillation batch size.** The colored outline indicates the input image. Increasing the batch size from 16 to 128 yields videos with superior visual quality and physical consistency. For instance, the spear in the rider's hand exhibits much greater temporal clarity. Additionally, the scaled model preserves exact object counts across frames (the eggs), whereas the smaller batch size struggles with physical plausibility by generating phantom objects out of nowhere (highlighted by the white circle).

uses it. As an additional control, we fine-tune the teacher for two epochs on exactly the same 30,000-video mixed-style set with DiffSynth-Studio ModelScope Team, and report it as Teacher (SFT) in Appendix Table 5. Standard supervised fine-tuning on these samples degrades both average quality and diversity relative to the original teacher, whereas routing the same samples through DFD improves the distilled student on both axes. This control shows that access to high-quality reference videos alone is not sufficient: the improvement depends on *how* DFD incorporates them, through the teacher-score discrepancy. DMD2 and DP-DMD in Table 1 are trained on the same reference data as DFD, so they remain the relevant controlled comparisons. More ablation study are in Appendix C

4.3 Scaling Up with Large Batch Size

We further scale up the image-to-video generation on the Cosmos-Predict2.5-2B model by increasing the batch size from 16 to 128. As shown quantitatively in Table 8 and visually in Fig 9, the model distilled with a large batch size consistently outperforms the model with a small batch size. The larger batch size yields remarkably better temporal stability during large motions (e.g., the rider's spear) and improved physical plausibility, such as preserving object permanence (e.g., consistent egg counts).

5 Conclusion

We introduce a data forcing distillation framework that advances the state-of-the-art for few-step video diffusion models, and restores diversity and fidelity. Our DFD requires only a single line of code change to the DMD2 baseline. At its core is the explicit integration of reference data into the distribution matching objective, which successfully overcomes the diversity degradation and over-saturation typically induced by reverse-KL formulations. Through comprehensive, large-scale evaluations on text-to-video, autoregressive video and image-to-video generation benchmarks, we demonstrate that our approach successfully resolves persistent failure modes in DMD, yielding substantial improvements in visual fidelity, temporal coherence, and physical plausibility in few-step video generation. We summarize the limitation of our work in Appendix F.4

6 AI USE STATEMENT

In this work, we used Qwen3-VL-4B-Instruct to annotate the video data. Apart from this Qwen-assisted annotation, we did not use generative AI tools to generate synthetic datasets; develop theoretical models or conceptual frameworks; formulate mathematical claims; provide critical ingredients for proving mathematical claims; assist in writing proofs; propose or refine hypotheses; design or provide feedback on research methodology or experiments; implement methods; assist with translation; otherwise clean or reformat datasets; support qualitative or thematic data analysis; or interpret results. The remaining required disclosure tasks are not applicable to this work.

Additionally, we used generative AI tools to create or modify scientific figures and images; draft parts of the paper; summarize and analyze existing literature; discover research topics and identify gaps; source and search for information; edit the paper to improve readability; identify relevant literature; format references; and suggest the paper's structure. AI-assisted data annotations were subject to our data-processing validation and human-in-the-loop selection procedures. The authors reviewed and edited all AI-assisted manuscript text and figures, and checked literature suggestions, citations, and factual claims against the original sources. We take responsibility for the final content of this work, including all text, claims, and artifacts produced with the aid of generative AI.

REFERENCES

Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Leo Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, et al. Video generation models as world simulators. *OpenAI Blog*, 1(8):1, 2024.

Boyuan Chen, Diego Martí Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. Diffusion forcing: Next-token prediction meets full-sequence diffusion. *Advances in Neural Information Processing Systems*, 37:24081–24125, 2024.

Zhengyang Geng, Ashwini Pokle, Weijian Luo, Justin Lin, and Zico Kolter. Consistency models made easy. In *International Conference on Learning Representations*, volume 2025, pp. 75158–75185, 2025.

Zhengyang Geng, Mingyang Deng, Xingjian Bai, Zico Kolter, and Kaiming He. Mean flows for one-step generative modeling. *Advances in Neural Information Processing Systems*, 38:75460–75482, 2026.

Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

Jiajun He, Wenlin Chen, Mingtian Zhang, David Barber, and José Miguel Hernández-Lobato. Training neural samplers with reverse diffusive kl divergence. *arXiv preprint arXiv:2410.12456*, 2024.

Jiahui Huang, Qunjie Zhou, Hesam Rabeti, Aleksandr Korovko, Huan Ling, Xuanchi Ren, Tianchang Shen, Jun Gao, Dmitry Slepichev, Chen-Hsuan Lin, Jiawei Ren, Kevin Xie, Joydeep Biswas, Laura Leal-Taixe, and Sanja Fidler. Vipe: Video pose engine for 3d geometric perception. In *NVIDIA Research Whitepapers*, 2025a.

Xun Huang, Zhengqi Li, Guande He, Mingyuan Zhou, and Eli Shechtman. Self forcing: Bridging the train-test gap in autoregressive video diffusion, 2025b. URL https://arxiv.org/abs/2506.08009.

Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, Yaohui Wang, Xinyuan Chen, Limin Wang, Dahua Lin, Yu Qiao, and Ziwei Liu. Vbench: Comprehensive benchmark suite for video generative models, 2023. URL https://arxiv.org/abs/2311.17982.

Dongjun Kim, Chieh-Hsin Lai, WeiHsiang Liao, Naoki Murata, Yuhta Takida, Toshimitsu Uesaka, Yutong He, Yuki Mitsufuji, and Stefano Ermon. Consistency trajectory models: Learning probability flow ode trajectory of diffusion. In *International Conference on Learning Representations*, volume 2024, pp. 44493–44525, 2024.

Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, Kathrina Wu, Qin Lin, Junkun Yuan, Yanxin Long, Aladdin Wang, Andong Wang, Changlin Li, Duojun Huang, Fang Yang, Hao Tan, Hongmei Wang, Jacob Song, Jiawang Bai, Jianbing Wu, Jinbao Xue, Joey Wang, Kai Wang, Mengyang Liu, Pengyu Li, Shuai Li, Weiyan Wang, Wenqing Yu, Xinchi Deng, Yang Li, Yi Chen, Yutao Cui, Yuanbo Peng, Zhentao Yu, Zhiyu He, Zhiyong Xu, Zixiang Zhou, Zunnan Xu, Yangyu Tao, Qinglin Lu, Songtao Liu, Dax Zhou, Hongfa Wang, Yong Yang, Di Wang, Yuhong Liu, Jie Jiang, and Caesar Zhong. Hunyuanvideo: A systematic framework for large video generative models, 2025. URL https://arxiv.org/abs/2412.03603.

Jiachen Li, Weixi Feng, Tsu-Jui Fu, Xinyi Wang, Sugato Basu, Wenhu Chen, and William Y Wang. T2v-turbo: Breaking the quality bottleneck of video consistency model with mixed reward feedback. *Advances in neural information processing systems*, 37:75692–75726, 2024.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling, 2023. URL https://arxiv.org/abs/2210.02747.

Cheng Lu and Yang Song. Simplifying, stabilizing and scaling continuous-time consistency models. In *International Conference on Learning Representations*, volume 2025, pp. 50611–50649, 2025.

Yanzuo Lu, Yuxi Ren, Xin Xia, Shanchuan Lin, Xing Wang, Xuefeng Xiao, Andy J Ma, Xiaohua Xie, and Jian-Huang Lai. Adversarial distribution matching for diffusion distillation towards efficient image and video synthesis. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 16818–16829. IEEE, 2025.

Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*, 2023a.

Weijian Luo, Tianyang Hu, Shifeng Zhang, Jiacheng Sun, Zhenguo Li, and Zhihua Zhang. Diff-instruct: A universal approach for transferring knowledge from pre-trained diffusion models. *Advances in Neural Information Processing Systems*, 36:76525–76546, 2023b.

Weijian Luo, Zemin Huang, Zhengyang Geng, J Zico Kolter, and Guo-jun Qi. One-step diffusion distillation through score implicit matching. *Advances in Neural Information Processing Systems*, 37:115377–115408, 2024.

Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and Tim Salimans. On distillation of guided diffusion models. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14297–14306. IEEE, 2023.

ModelScope Team. Diffsynth-studio. https://github.com/modelscope/DiffSynth-Studio. GitHub repository.

Weili Nie, Julius Berner, Nanye Ma, Chao Liu, Saining Xie, and Arash Vahdat. Transition matching distillation for fast video generation. *arXiv preprint arXiv:2601.09881*, 2026.

NVIDIA, :, Arslan Ali, Junjie Bai, Maciej Bala, Yogesh Balaji, Aaron Blakeman, Tiffany Cai, Jiaxin Cao, Tianshi Cao, Elizabeth Cha, Yu-Wei Chao, Prithvijit Chattopadhyay, Mike Chen, Yongxin Chen, Yu Chen, Shuai Cheng, Yin Cui, Jenna Diamond, Yifan Ding, Jiaojiao Fan, Linxi Fan, Liang Feng, Francesco Ferroni, Sanja Fidler, Xiao Fu, Ruiyuan Gao, Yunhao Ge, Jinwei Gu, Aryaman Gupta, Siddharth Gururani, Imad El Hanafi, Ali Hassani, Zekun Hao, Jacob Huffman, Joel Jang, Pooya Jannaty, Jan Kautz, Grace Lam, Xuan Li, Zhaoshuo Li, Maosheng Liao, Chen-Hsuan Lin, Tsung-Yi Lin, Yen-Chen Lin, Huan Ling, Ming-Yu Liu, Xian Liu, Yifan Lu, Alice Luo, Qianli Ma, Hanzi Mao, Kaichun Mo, Seungjun Nah, Yashraj Narang, Abhijeet Panaskar, Lindsey Pavao, Trung Pham, Morteza Ramezanali, Fitsum Reda, Scott Reed, Xuanchi Ren, Haonan Shao, Yue Shen, Stella Shi, Shuran Song, Bartosz Stefaniak, Shangkun Sun, Shitao Tang, Sameena Tasmeen, Lyne Tchapmi, Wei-Cheng Tseng, Jibin Varghese, Andrew Z. Wang, Hao Wang, Haoxiang Wang, Heng Wang, Ting-Chun Wang, Fangyin Wei, Jiashu Xu, Dinghao Yang, Xiaodong Yang, Haotian Ye, Seonghyeon Ye, Xiaohui Zeng, Jing Zhang, Qinsheng Zhang, Kaiwen Zheng, Andrew Zhu, and Yuke Zhu. World simulation with video foundation models for physical ai, 2026. URL https://arxiv.org/abs/2511.00062.

Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion, 2022. URL https://arxiv.org/abs/2209.14988.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PmLR, 2021.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2022. URL https://arxiv.org/abs/2112.10752.

Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align your flow: Scaling continuous-time flow map distillation, 2025. *URL https://arxiv. org/abs/2506.14603*.

Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022.

Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion distillation. In *European Conference on Computer Vision*, pp. 87–103. Springer, 2024.

Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. Dinov3. *arXiv preprint arXiv:2508.10104*, 2025.

Yang Song and Prafulla Dhariwal. Improved techniques for training consistency models. In *International Conference on Learning Representations*, volume 2024, pp. 15078–15097, 2024.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations, 2021. URL https://arxiv.org/abs/2011.13456.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023.

Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wente Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. Wan: Open and advanced large-scale video generative models, 2025. URL https://arxiv.org/abs/2503.20314.

Fu-Yun Wang, Zhaoyang Huang, Alexander W Bergman, Dazhong Shen, Peng Gao, Michael Lingelbach, Keqiang Sun, Weikang Bian, Guanglu Song, Yu Liu, et al. Phased consistency models. *Advances in neural information processing systems*, 37:83951–84009, 2024.

Yifei Wang, Weimin Bai, Debing Zhang, Weijian Luo, He Sun, et al. Uni-instruct: One-step diffusion model through unified diffusion divergence instruction. *Advances in Neural Information Processing Systems*, 38:57097–57132, 2026.

Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation, 2023. URL https://arxiv.org/abs/2305.16213.

Tianhe Wu, Ruibin Li, Lei Zhang, and Kede Ma. Diversity-preserved distribution matching distillation for fast visual synthesis, 2026. URL https://arxiv.org/abs/2602.03139.

Jiazheng Xu, Yu Huang, Jiale Cheng, Yuanming Yang, Jiajun Xu, Yuan Wang, Wenbo Duan, Shen Yang, Qunlin Jin, Shurun Li, et al. Visionreward: Fine-grained multi-dimensional human preference learning for image and video generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 11269–11277, 2026.

Yilun Xu, Weili Nie, and Arash Vahdat. One-step diffusion models with f-divergence distribution matching. *arXiv preprint arXiv:2502.15681*, 2025.

Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, Da Yin, Yuxuan Zhang, Weihan Wang, Yean Cheng, Bin Xu, Xiaotao Gu, Yuxiao Dong, and Jie Tang. Cogvideox: Text-to-video diffusion models with an expert transformer, 2025. URL https://arxiv.org/abs/2408.06072.

Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6613–6623. IEEE, 2024a.

Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and William T. Freeman. Improved distribution matching distillation for fast image synthesis, 2024b. URL https://arxiv.org/abs/2405.14867.

Tianwei Yin, Qiang Zhang, Richard Zhang, William T Freeman, Fredo Durand, Eli Shechtman, and Xun Huang. From slow bidirectional to fast autoregressive video diffusion models. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 22963–22974. IEEE, 2025.

Kaiwen Zheng, Yuji Wang, Qianli Ma, Huayu Chen, Jintao Zhang, Yogesh Balaji, Jianfei Chen, Ming-Yu Liu, Jun Zhu, and Qinsheng Zhang. Large scale diffusion distillation via score-regularized continuous-time consistency. In *International Conference on Learning Representations*, volume 2026, pp. 2582–2603, 2026.

Mingyuan Zhou, Huangjie Zheng, Zhendong Wang, Mingzhang Yin, and Hai Huang. Score identity distillation: Exponentially fast distillation of pretrained diffusion models for one-step generation. In *Forty-first International Conference on Machine Learning*, 2024.

Hongzhou Zhu, Min Zhao, Guande He, Hang Su, Chongxuan Li, and Jun Zhu. Causal forcing: Autoregressive diffusion distillation done right for high-quality real-time interactive video generation. *arXiv preprint arXiv:2602.02214*, 2026.

A Related Work

A.1 Trajectory-Based Methods

Trajectory-based distillation methods aim to accelerate the sampling process by training a few-step student model to match the denoising trajectory of a multi-step teacher model Song et al. (2023); Luo et al. (2023a); Meng et al. (2023); Li et al. (2024). Among these, Consistency Models (CMs) Song et al. (2023); Song & Dhariwal (2024); Geng et al. (2025; 2026); Sabour et al.; Kim et al. (2024) learns a consistency function $f_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t) \to \boldsymbol{x}_0$ that directly maps any point $\boldsymbol{x}_t$ on the teacher's probability flow ODE trajectory to the initial data point $\boldsymbol{x}_0$. Initially, discrete-time CMs were trained by minimizing the discrepancy between outputs at adjacent anchor points t and $t - \Delta t$ along the teacher trajectory. Recently, continuous-time CMs Lu & Song (2025); Kim et al. (2024) offer a clean upgrade by taking the limit as $\Delta t \to 0$, simplifying the objective to enforce instantaneous self-consistency. Building upon these foundations, a notable work is the rCM Zheng et al. (2026), which is the first work to scale up continuous-time consistency models to large image and video diffusion models. Despite their significant empirical success and clean mathematical formulations, trajectory-based distillation methods still face open challenges, often exhibiting noticeable performance degradation when scaled up to large pretrained image synthesis or video generation models Wang et al. (2024).

A.2 Distribution-Matching Methods

Distribution-matching distillation compresses a multi-step teacher into a few-step student by aligning the student's output distribution with the teacher's. DMD Yin et al. (2024a) minimizes a reverse KL between the two and achieves stable training and strong sample quality; DMD2 Yin et al. (2024b) augments it with an auxiliary GAN discriminator to mitigate the mode-seeking behavior of reverse KL, but inherits the well-known instability of adversarial training and only injects the reference-data signal *implicitly* through the discriminator. These issues are amplified for video diffusion, where the high spatiotemporal dimensionality makes mode collapse and over-saturation considerably more severe.

A parallel line of work generalizes the divergence itself. f-distill Xu et al. (2025) casts DMD as a special case of integral f-divergence and reports gains from alternatives such as the Jeffreys divergence, while Uni-Instruct Wang et al. (2026) unifies Score Implicit Matching Luo et al. (2024) and Score Identity Distillation Zhou et al. (2024) under a single objective. Diversity-Preserving DMD Wu et al. (2026) instead couples DMD with trajectory-based distillation to recover modal diversity, though scaling such hybrids to video remains open. Most recently, Transition Matching Distillation (TMD) Nie et al. (2026) extends the paradigm temporally by matching conditional transition probabilities across sub-intervals of the denoising trajectory, enabling distillation of large-scale video models – yet it still rests on the same reverse-KL backbone and inherits its diversity limitations.

Distribution Matching for Autoregressive Video Generation. The same machinery drives the recent progress in autoregressive video generation, where a video is produced chunk by chunk under causal attention so that long-horizon and streaming synthesis become possible. Diffusion Forcing Chen et al. (2024) trains sequence diffusion models with independent per-token noise levels, combining next-token prediction with full-sequence diffusion. CausVid Yin et al. (2025) then converts a slow bidirectional teacher into a fast causal few-step student through distribution matching distillation, and Self Forcing Huang et al. (2025b) closes the remaining train-test gap by rolling out the student's own predictions during training under the DMD objective, while for the Causal Forcing Zhu et al. (2026), it refines this recipe for high-quality real-time interactive generation. Since these causal students are optimized with the same reverse-KL gradient as their bidirectional counterparts, the mode-seeking behavior discussed above carries over unchanged, and DFD applies to them without modification (Sec. 4.1.3).

B DFD PSEUDOCODE

Algorithm 1 highlights the simplicity of DFD, which requires only a *one-line change* to the original DMD2 student update.

Algorithm 1 Pseudocode of Student Update Step in a PyTorch-like Style.

```
# student_network / teacher_network / fake_score_network: networks
# input_student, t_student: noisy input and timestep fed to student
# t, eps: diffusion timestep and noise for forward diffusion
# data: reference data batch; condition: video gen conditioning
# p: probability of using DFD vs DMD
def student_update_step(input_student, t_student, t, eps, data, condition
    =None):
    # Generate samples using student network.
    gen_data = student_network(input_student, t_student, condition=
    condition)

    # teacher_data = gen_data.detach() # Original DMD update
    teacher_data = data.detach() if (torch.rand() < p) else
    gen_data.detach()

    # Inject noise into the data via forward diffusion
    perturbed_data         = forward_diffusion(gen_data, eps, t)
    perturbed_teacher_data = forward_diffusion(teacher_data, eps, t)
    # Estimate scores
    fake_score    = fake_score_network(perturbed_data, t, condition=
    condition)
    teacher_score = teacher_network(perturbed_teacher_data, t, condition=
    condition)
    # Compute gradient and student loss
    vsd_grad      = fake_score - teacher_score
    pseudo_target = gen_data - vsd_grad
    gen_loss      = 0.5 * F.mse_loss(gen_data, pseudo_target.detach())
    return gen_loss
```

The violet block is the *only* difference between the DFD and DMD frameworks.

C ADDITIONAL RESULTS FOR ABLATION STUDY

Table 5: Reference-data control for text-to-video generation. The teacher and its supervised fine-tuned variant use the same evaluation protocol as Table 1. For all metrics, higher is better.

	Video Diversity		Camera Pose Diversity		Video Quality					
Method	CLIP (Per-frame)	DINO (Per-frame)	Endpoint Distance	Trajectory Distance	Subject Consistency	Background Consistency	Temporal Flickering	Motion Smoothness	Aesthetic Quality	Average VBench
Teacher	**0.222**	**0.350**	**30.571**	**26.651**	**0.956**	**0.959**	**0.976**	**0.985**	0.622	**0.899**
Teacher (SFT)	0.186	0.295	9.977	9.694	0.935	0.951	0.973	0.984	**0.623**	0.893

Effects of the GAN Loss. We ablate the requirement of GAN loss, and compare our DFD pipeline with and without an additional GAN loss. We experiment on the text-to-video generation on animation set. The VBench quality scores are reported in Table 6, and qualitative results in Fig. 6(a). This ablation suggests that, in our post-training setting, adding the GAN loss does not provide a consistent benefit. Removing it simplifies the pipeline and improves dynamic degree and imaging quality, while maintaining comparable scores on the remaining metrics.

Effects on Weights in Eq. 9. We ablate the weights in Eq. 9. We experiment on text-to-video generation using the animation set. We compare two weighting schemes, $w = \frac{1}{2}$ (our default setting) and $w = 1$ (pure DFD without DMD). Quantitative results are reported in Table 7, and qualitative results in Fig. 6(b). The differences between the two settings are small, indicating that our method



Figure 6: **(a): Ablation on the effect of the GAN loss.** Adding the GAN loss yields no clear quality improvement, and video dynamics even degrade, which is consistent with our quantitative results. **(b): Ablation on the weight in Eq. 9**: $w = 1$ (upper) vs. $w = \frac{1}{2}$ (lower). Each row shows three evenly sampled frames from one generated video. There is no clear visual difference between the two settings. **(c): Qualitative results without DMD2 pretraining.** Initializing purely from the teacher model without DMD2 pretraining fails to produce reasonable results, even after a sufficient number of training steps for both the text-to-video and image-to-video tasks.

Table 6: Ablation on the effects of GAN loss in our DFD. Bold indicates the better value. Removing the GAN loss yields results comparable to the model distilled with the GAN loss.

Model (DFD)	Subject Consistency	Background Consistency	Temporal Flickering	Motion Smoothness	Dynamic Degree	Aesthetic Quality	Imaging Quality
w/o GAN	**0.9690**	0.9620	0.9785	0.9899	**0.5000**	0.7194	**0.7452**
w GAN	0.9666	**0.9625**	**0.9831**	**0.9912**	0.3750	**0.7213**	0.7210

is stable and insensitive to the choice of w. We adopt $w = \frac{1}{2}$ for theoretical reasons, as it better satisfies the condition in Eq. 6.

Table 7: Ablation on the weights in Eq. 9. The results show no clear difference between the two choices of w.

Model	Subject Consistency	Background Consistency	Temporal Flickering	Motion Smoothness	Dynamic Degree	Aesthetic Quality	Imaging Quality
$w = \frac{1}{2}$	**0.9691**	**0.9661**	0.9830	0.9907	**0.5625**	0.7135	**0.7457**
$w = 1$	0.9666	0.9625	**0.9831**	**0.9912**	0.5450	**0.7213**	0.7210

Effects of DMD2 Pretraining. Our DFD needs to satisfy the *validity condition* in Eq. 6 for stabilized training. To verify this, we ablate by initializing the student with the teacher model (without DMD2 distillation) and comparing the DMD2 initialization. We experiment on both text-to-video generation with the animation set and image-to-video generation with the mix-style set. Qualitative

results are provided in Fig. 6(c). Even after a long training run (e.g., 1400 iterations), the model fails to converge, which underscores the necessity of the validity condition.

D Additional Theory

D.1 Teacher score discrepancy

We show the expectation of the teacher score discrepancy equals to zero when the student is perfect:

Proposition. If the student matches the data distribution exactly ($p_{\text{fake}} = p_{\text{refer}}$), then $\mathbb{E}_{\boldsymbol{x},\boldsymbol{z}}[\Delta_{p_{\text{refer}},\boldsymbol{x},G_\theta(\boldsymbol{z})}] = 0$.

Proof. When the student is perfect, $\boldsymbol{x} = G_\theta(\boldsymbol{z})$ is a sample from p_{refer}, and $\hat{\boldsymbol{x}}$ is also a sample from p_{refer}. They are i.i.d. and therefore *exchangeable*. With shared $\boldsymbol{\epsilon}$,

$$\boldsymbol{x}_t^{\text{fake}} = \alpha_t \boldsymbol{x} + \sigma_t \boldsymbol{\epsilon} \quad \stackrel{d}{=} \quad \alpha_t \hat{\boldsymbol{x}} + \sigma_t \boldsymbol{\epsilon} = \boldsymbol{x}_t^{\text{refer}}. \tag{11}$$

Both have the same distribution $p_{\text{refer},t}$, so

$$\mathbb{E}_{\boldsymbol{x},\boldsymbol{z}}[\Delta_{p_{\text{refer}},\boldsymbol{x},G_\theta(\boldsymbol{z})}] = \mathbb{E}_{\boldsymbol{x},\boldsymbol{z}}[\nabla_{\boldsymbol{x}} \log p_{\text{refer}}(\boldsymbol{x}_t^{\text{refer}})] - \mathbb{E}_{\boldsymbol{x},\boldsymbol{z}}[\nabla_{\boldsymbol{x}} \log p_{\text{refer}}(\boldsymbol{x}_t^{\text{fake}}))] = 0. \tag{12}$$

$\Delta_{p_{\text{refer}},\boldsymbol{x},G_\theta(\boldsymbol{z})}$ is a zero-mean stochastic perturbation at the optimum. □

D.2 Bounded MSE between $g_{\text{DFD}}(\theta)$ and $\nabla_\theta D_{\text{KL}}(p_{\text{fake}} \| p_{\text{refer}})$

Setup

In score-distillation methods, the teacher score is never evaluated on clean inputs. Following standard practice, samples are perturbed by the diffusion forward process to a noise level $t \in (0, T]$ before being passed to the teacher. Let

$$\boldsymbol{x}_t := \alpha_t \boldsymbol{x} + \sigma_t \boldsymbol{\epsilon}, \boldsymbol{x} \sim p_{\text{refer}}, \qquad \tilde{\boldsymbol{x}}_t := \alpha_t G_\theta(\boldsymbol{z}, c) + \sigma_t \boldsymbol{\epsilon}, \qquad \boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}), \tag{13}$$

where (α_t, σ_t) are the diffusion schedule coefficients. The relevant scores are those of the noised marginals,

$$s_{\text{refer}}(\boldsymbol{x}_t \mid c, t) := \nabla_{\boldsymbol{x}_t} \log p_{\text{refer},t}(\boldsymbol{x}_t \mid c), \qquad s_{\text{fake}}(\tilde{\boldsymbol{x}}_t \mid c, t) := \nabla_{\tilde{\boldsymbol{x}}_t} \log p_{\text{fake},t}(\tilde{\boldsymbol{x}}_t \mid c), \tag{14}$$

and the teacher is a neural network $s_\phi(\cdot, c, t)$ approximating $s_{\text{refer}}(\cdot \mid c, t)$.

The DFD and KL gradient at noise level t are:

$$g_{\text{DFD}}(\theta) = \mathbb{E}_{\boldsymbol{x},\boldsymbol{z},\boldsymbol{\epsilon}}\big[\big(s_{\text{fake}}(\tilde{\boldsymbol{x}}_t \mid c, t) - s_{\text{refer}}(\boldsymbol{x}_t \mid c, t)\big)\nabla_\theta G_\theta(\boldsymbol{z}, c)\big], \tag{15}$$

$$\nabla_\theta D_{\text{KL}} = \mathbb{E}_{\boldsymbol{z},\boldsymbol{\epsilon}}\big[\big(s_{\text{fake}}(\tilde{\boldsymbol{x}}_t \mid c, t) - s_{\text{refer}}(\tilde{\boldsymbol{x}}_t \mid c, t)\big)\nabla_\theta G_\theta(\boldsymbol{z}, c)\big]. \tag{16}$$

(For brevity we have absorbed any α_t chain-rule factors into $\nabla_\theta G_\theta$ without changing the structure of the argument.)

Condition

Condition 1 (Bounded conditional matching error)**.** There exists $\delta(c) \geq 0$ such that

$$\mathbb{E}\Big[\|\boldsymbol{x} - G_\theta(\boldsymbol{z}, c)\|^2 \;\Big|\; c\Big] \leq \delta(c)^2. \tag{17}$$

Condition 2 (Lipschitz teacher score on the noised diffusion path)**.** At each noise level $t > 0$, the teacher score network $s_{\text{refer}}(\cdot \mid c, t)$ is $L(c, t)$-Lipschitz:

$$\|s_{\text{refer}}(\boldsymbol{u} \mid c, t) - s_{\text{refer}}(\boldsymbol{v} \mid c, t)\| \leq L(c, t)\, \|\boldsymbol{u} - \boldsymbol{v}\|, \qquad \forall \boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^d. \tag{18}$$

This is justified for two complementary reasons: (i) the noised marginal $p_{\mathrm{refer},t}$ is the convolution of p_{refer} with a Gaussian of variance σ_t^2, whose score has Hessian operator norm bounded by $\mathcal{O}(1/\sigma_t^2)$; (ii) the score is parameterized by a neural network with Lipschitz architectural primitives, so it inherits a finite Lipschitz constant on any bounded domain.

Assumption 1 (Bounded generator Jacobian)**.** Let $\boldsymbol{z} \sim \mathcal{N}(0, I)$ denote the latent noise and let c denote the conditioning variable. We assume that there exists a finite constant $B > 0$ such that

$$\mathbb{E}_{\boldsymbol{z}}[\|\nabla_\theta G_\theta(\boldsymbol{z}, c)\|_{\mathrm{op}}^2 \mid c] \le B^2 < \infty.$$

The Bounded-generator-Jacobian assumption is mild for modern diffusion-based video generators. In architectures such as Wan, the generator is a finite-depth diffusion transformer composed of linear projections, attention blocks, normalization layers, and smooth nonlinearities such as SiLU, while the latent input is sampled from a Gaussian distribution. For such networks, the parameter Jacobian is expected to grow at most polynomially with the latent magnitude under finite model parameters. Since Gaussian random variables possess finite moments of all orders, polynomial growth of the Jacobian is sufficient to guarantee

$$\mathbb{E}_{\boldsymbol{z}}[\|\nabla_\theta G_\theta(\boldsymbol{z}, c)\|_{\mathrm{op}}^2 \mid c] \le B^2 < \infty.$$

Importantly, this condition does not require the Jacobian to be uniformly bounded for every possible latent sample; it only requires its squared magnitude to be integrable under the latent distribution actually encountered during generation.

MAIN RESULT

Proposition 1. Under Condition 1–2 and assumption 1, for every noise level $t > 0$,

$$\left\|g_{\mathrm{DFD}}(\theta) - \nabla_\theta D_{\mathrm{KL}}(p_{\mathrm{fake}} \| p_{\mathrm{refer}})\right\|_2^2 \;\le\; B^2\, \alpha_t^2\, L(c,t)^2\, \delta(c)^2. \tag{19}$$

Proof. **Step 1: Cancellation of the fake score.** The fake-score term in (15) depends only on $(\boldsymbol{z}, \boldsymbol{\epsilon}, c)$, not on $\boldsymbol{x}$, so $\mathbb{E}_{\boldsymbol{x},\boldsymbol{z},\boldsymbol{\epsilon}|c}[\,\cdot\,] = \mathbb{E}_{\boldsymbol{z},\boldsymbol{\epsilon}|c}[\,\cdot\,]$ for that term. It cancels with the corresponding term in (16), giving

$$g_{\mathrm{DFD}}(\theta) - \nabla_\theta D_{\mathrm{KL}} = \mathbb{E}_{\boldsymbol{x},\boldsymbol{z},\boldsymbol{\epsilon}|c}\big[\big(s_{\mathrm{refer}}(\tilde{\boldsymbol{x}}_t \mid c, t) - s_{\mathrm{refer}}(\boldsymbol{x}_t \mid c, t)\big)\nabla_\theta G_\theta(\boldsymbol{z}, c)\big]\,. \tag{20}$$

Step 2: Jensen's inequality. Since $\|\cdot\|_2^2$ is convex,

$$\left\|g_{\mathrm{DFD}}(\theta) - \nabla_\theta D_{\mathrm{KL}}\right\|_2^2 \;\le\; \mathbb{E}_{\boldsymbol{x},\boldsymbol{z},\boldsymbol{\epsilon}|c}\Big[\left\|\big(s_{\mathrm{refer}}(\tilde{\boldsymbol{x}}_t \mid c, t) - s_{\mathrm{refer}}(\boldsymbol{x}_t \mid c, t)\big)\nabla_\theta G_\theta(\boldsymbol{z}, c)\right\|_2^2\Big]\,. \tag{21}$$

Step 3: Bounding the integrand. Using sub-multiplicativity, Condition 2, and Assumption 1,

$$\begin{aligned}\left\|\big(s_{\mathrm{refer}}(\tilde{\boldsymbol{x}}_t \mid c, t) - s_{\mathrm{refer}}(\boldsymbol{x}_t \mid c, t)\big)\nabla_\theta G_\theta(\boldsymbol{z}, c)\right\|_2^2 &\le B^2\, \left\|s_{\mathrm{refer}}(\tilde{\boldsymbol{x}}_t \mid c, t) - s_{\mathrm{refer}}(\boldsymbol{x}_t \mid c, t)\right\|^2 \\ &\le B^2\, L(c,t)^2\, \left\|\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_t\right\|^2.\end{aligned} \tag{22}$$

Since $\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_t = \alpha_t\big(G_\theta(\boldsymbol{z}, c) - \boldsymbol{x}\big)$ (the noise term $\sigma_t\boldsymbol{\epsilon}$ is shared and cancels),

$$\left\|\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_t\right\|^2 = \alpha_t^2\, \left\|G_\theta(\boldsymbol{z}, c) - \boldsymbol{x}\right\|^2. \tag{23}$$

Step 4: Apply the matching bound. Combining the previous steps,

$$\begin{aligned}\left\|g_{\mathrm{DFD}}(\theta) - \nabla_\theta D_{\mathrm{KL}}\right\|_2^2 &\le B^2\, \alpha_t^2\, L(c,t)^2\, \mathbb{E}\Big[\left\|\boldsymbol{x} - G_\theta(\boldsymbol{z}, c)\right\|^2 \;\Big|\; c\Big] \\ &\le B^2\, \alpha_t^2\, L(c,t)^2\, \delta(c)^2. \quad \square\end{aligned} \tag{24}$$

Table 8: Comparison between models distilled with batch size 16 and 128, evaluated on the VBench test set.

Batch size	Subject Consistency	Background Consistency	Aesthetic Quality	Temporal Flickering	Motion Smoothness	Image-to-Video Subject	Image-to-Video Background	Average
16	0.9613	**0.9692**	0.6371	**0.9759**	0.9900	0.9859	**0.9930**	0.9303
128	**0.9638**	0.9685	**0.6383**	0.9783	**0.9914**	**0.9878**	0.9929	**0.9316**

Additional discussion: Alignment Between the Reference Sample and the Generated Sample

This distinction between the reference sample and the generated sample is precisely the motivation for the form of the validity condition in Eq. (6): it is an expectation-level condition on the conditional reference and generated distributions, and it does *not* require every sampled pair to be close. The distinction matters in both of our settings. In text-to-video, two videos sharing the same text prompt may differ substantially in semantic content, yet DFD is concerned with matching their conditional distributions rather than with establishing an instance-level correspondence. Image-to-video adds a first-frame condition, but the future evolution remains multimodal; the additional conditioning therefore narrows the conditional distribution without turning the reference and generated videos into a one-to-one matched pair.

Under this distribution-level condition, together with a Lipschitz teacher score and a bounded generator Jacobian (Conditions 1–1 in Appendix D), Proposition 1 bounds the discrepancy between the DFD gradient and the original reverse-KL gradient,

$$\left\| g_{\mathrm{DFD}}(\theta) - \nabla_\theta D_{\mathrm{KL}}\left(p_{\mathrm{fake}} \,\|\, p_{\mathrm{reference}}\right) \right\|_2^2 \;\le\; B^2\, \alpha_t^2\, L(c,t)^2\, \delta(c)^2. \tag{25}$$

Thus our theoretical claim is not that an arbitrary reference sample always supplies a correct instance-level direction. Rather, DFD is a *controlled approximation* to the original reverse-KL gradient whenever the two conditional distributions are already sufficiently close in the sense of Eq. (6): the two gradients coincide in the limiting case $\delta(c) = 0$, and for nonzero $\delta(c)$ Eq. (25) gives an explicit approximation-error bound. Together with the fact that the teacher-score discrepancy vanishes in expectation when $p_{\mathrm{fake}} = p_{\mathrm{reference}}$.

E Experiment Details

E.1 Data Curation

We mainly curated two datasets for our experiments:

- *animation dataset*: video clips with an animation or cartoon visual style.
- *mix-style dataset*: video clips spanning all visual styles.

Both are derived from the ViPE dataset Huang et al. (2025a), starting from ∼966k captioned source clips. Our pipeline has five stages: (1) tri-level captioning, (2) CLIP visual feature extraction, (3) K-means clustering, (4) per-cluster grid visualization with human-in-the-loop selection, and (5) WebDataset packaging at three caption granularities. The full pipeline code is released alongside the paper.

E.1.1 Tri-Level Captioning.

For each source video we generate three captions of different lengths in a single VLM forward pass, using **Qwen3-VL-4B-Instruct** served via vLLM with continuous batching and PagedAttention. A single system prompt instructs the model to emit

- `[LONG]`: a 250–300 word caption covering subject appearance, actions, expressions, environment, camera work, and visual style;
- `[MEDIUM]`: a 150–180 word caption that retains the key subject, actions, setting, and style;

- `[SHORT]`: a single sentence capturing the essence of the clip.

Captions are decoded with sampling parameters $T = 0.7$, top-$p = 0.8$, and `max_new_tokens`$=512$. Producing all three lengths in one forward pass (rather than three independent passes) reduces VLM cost roughly $3\times$ since the visual tokens are encoded only once. The three captions are parsed from the tagged output via regular expressions; if any tag is missing we fall back deterministically (truncate the long caption for medium, take the first sentence for short). Each caption is also encoded with the UMT5-XXL text encoder (Wan2.1-T2V-1.3B checkpoint) and the resulting prompt embeddings are stored alongside the raw text.

E.1.2 Visual Feature Extraction.

For diversity sampling we want a feature space that reflects visual content rather than caption phrasing. We therefore extract CLIP ViT-B/32 image feature. For each video we decode 3 uniformly-spaced keyframes with PyAV, apply the standard CLIP preprocessing (`Resize`(224), `CenterCrop`(224), `ToTensor`, CLIP normalization), encode each keyframe with `encode_image`, mean-pool the three frame embeddings, and ℓ_2-normalize, yielding one 512-d vector per video.

E.1.3 Clustering.

We cluster the $\sim$966k feature vectors with **FAISS** K**-means** (K=1000, 50 iterations). Because all vectors are ℓ_2-normalized, squared Euclidean distance to the centroid is monotonically related to cosine similarity, so we use it directly for ranking inside each cluster. The output is a per-video integer cluster label and a $(K,512)$ centroid matrix.

E.1.4 Per-Cluster Visualization and Human-in-the-Loop Selection.

For each of the K=1000 clusters we (i) sort all members by Euclidean distance from the centroid, (ii) pick 80 candidates at evenly-spaced percentiles of that distance (so row 0 is the centroid-nearest video and row 39 is the farthest), and (iii) decode 5 evenly-spaced frames per candidate. A human reviewer then opens each cluster grid and writes the row indices of 40 visually-representative-but-distinct videos into the `picks` column, e.g. `"0,2,5,7,9,12,...,68,79"`. It is possible that a whole cluster is completely dropped.The even-percentile sampling ensures the candidates already span the full intra-cluster spread, so the reviewer is choosing among visually distinct options rather than near-duplicates. Clusters whose grid contains no useful content (e.g., uniformly black frames, decode-failed clusters) can be left empty and are dropped from the final dataset; this trades a small number of slots for higher quality. After review we obtain 30,000 selected videos in cluster-id order.

The animation dataset is constructed by the same filtering pipeline, restricted to clusters whose grid visualizations are dominated by animation/cartoon style during the human review step.

E.2 Baseline Details

We show the implementation details of our baselines in Table 9. For the DP-DMD Wu et al. (2026), it doesn't need the GAN in the training process. For the text-to-video tasks and autoregressive generation tasks, all the baseline and our method are trained on 8 A100 GPUs. For the image-to-video tasks, all the baseline and our method are trained on 16 A100 GPUs.

Table 9: Training configurations.

Models	Cosmos Predict2 I2V			Wan2.1 T2V			Wan2.1 Autoregressive		
	DMD2	DP-DMD	DFD	DMD2	DP-DMD	DFD	Causal Forcing	Self Foring(DMD)	Self Forcing(DFD)
Num. of Frame	81	81	81	81	81	81	81	81	81
Batch Size	16	16	16	16	16	16	16	16	16
Resolution	480p	480p	480p	480p	480p	480p	480p	480p	480p
Learning Rate(discriminator)	1e-05	N/A	N/A	1e-05	N/A	N/A	N/A	N/A	N/A
Learning Rate (student)	1e-05	1e-05	1e-05	1e-05	1e-05	1e-05	2e-06	2e-06	2e-06
Learning Rate (fake score)	1e-05	1e-05	1e-05	1e-05	1e-05	1e-05	4e-07	4e-07	4e-07
CFG Scale	3.0	3.0	3.0	5.0	5.0	5.0	3.0	3.0	3.0
Student Update Frequency	5	5	5	5	5	5	5	5	5
Diversity anchor step	N/A	1	N/A	N/A	5	N/A	N/A	N/A	N/A
Diversity weights	N/A	0.05	N/A	N/A	0.05	N/A	N/A	N/A	N/A
Total Iterations	3k	3k	300	3k	3k	100	1.2k	1.2k	70
Pretrained Model Iterations	0	0	2.5k(from DMD2)	0	0	2.5k(from DMD2)	0	0	1.1k(from Self-Forcing(DMD))

E.3 EVALUTION METRICS

Metric Details We show the metric details here. For each text prompt we sample $L = 8$ videos from distinct initial noise seeds, and report three families of metrics: the visual diversity metrics adapted from DP-DMD, a set of VBench Huang et al. (2023) quality dimensions, the camera-pose statistics extracted with ViPE Huang et al. (2025a) and the VisionReward metricsXu et al. (2026) following Causal Forcing Zhu et al. (2026)

E.3.1 VISUAL DIVERSITY METRIC.

We adopt the diversity metric from DP-DMD Wu et al. (2026). We assess sample diversity using DINOv2-ViT-Large (DINO) Siméoni et al. (2025) and CLIP-ViT-Large (CLIP) Radford et al. (2021) by computing the cosine similarity between extracted image feature representations in a pairwise manner:

$$\text{Diversity} = 1 - \frac{2}{L(L-1)} \sum_{i,\,j} \cos\left(\boldsymbol{x}_\theta^{(i)}, \boldsymbol{x}_\theta^{(j)}\right), \tag{26}$$

where L denotes the number of distinct initial noise samples per text prompt[1], and we set $L = 8$ in our experiments. To adapt this image-level metric to videos, we sample 16 evenly-spaced frames from each generated clip, extract per-frame CLIP and DINO embeddings, and apply Eq. 26 and average over all these 16 frames. The reported diversity is the average over all text prompts.

E.3.2 VBENCH QUALITY METRICS.

For t2v tasks, We evaluate generation quality with VBench Huang et al. (2023) which scores videos directly without requiring access to the original training prompts. We calculate the seven prompt-free dimensions and report the average VBench score:

- *subject consistency*: measures whether the main subject's appearance (identity, shape, texture) remains stable across frames, computed via DINO feature similarity between frames.
- *background consistency*: measures temporal stability of the background scene across frames using CLIP feature similarity, capturing flicker or drift in the surrounding environment.
- *temporal flickering*: penalizes high-frequency, low-level fluctuations between adjacent frames in static or near-static regions, reflecting visible flicker artifacts.
- *motion smoothness*: assesses whether motion follows the priors of a video frame interpolation model, with smoother (more physically plausible) motion scoring higher.
- *dynamic degree*: estimates the amount of motion in the video using optical flow magnitude, rewarding generations that are not nearly-static.
- *aesthetic quality*: predicts the human-perceived aesthetic score of sampled frames using the LAION aesthetic predictor.
- *imaging quality*: predicts low-level image quality (sharpness, noise, distortion) using the MUSIQ image-quality predictor.

For i2v tasks, we evaluate the following metrics provided in VBench i2v part and report the average score:

- *subject consistency:* Measures the model's ability to preserve the identity, appearance, and structural integrity of the primary subject across all frames of the generated video.
- *Background Consistency:* Evaluates the temporal stability and visual coherence of the background environment, penalizing unnatural shifts or morphing over time.
- *Aesthetic Quality:* Assesses the overall visual appeal, lighting, composition, and artistic fidelity of the individual generated frames.
- *Temporal Flickering:* Quantifies the presence of high-frequency visual artifacts or unnatural, rapid changes in texture and lighting between adjacent frames.

[1]For each text prompt, every pair of generated samples is compared once, yielding $\binom{L}{2}$ total comparisons.

- *Motion Smoothness:* Evaluates the naturalness, fluidity, and physical plausibility of the movements occurring within the video sequence.
- *Image-to-Video Subject:* Specifically for Image-to-Video (I2V) generation tasks, this metric calculates how accurately the subject in the generated video aligns with the reference subject provided in the conditioning input image.
- *Image-to-Video Background:* Evaluates how well the generated video maintains and reflects the background elements present in the initial conditioning image.

E.3.3 CAMERA-POSE DIVERSITY.

To quantify how the camera moves across generated videos, we run ViPE Huang et al. (2025a) on every clip to recover an 81-frame sequence of 4×4 camera-to-world transformation matrices $\{T_t\}_{t=0}^{80}$, with translation $\boldsymbol{p}_t = T_t[:3, 3]$ and rotation $R_t = T_t[:3, :3]$. We then derive two groups of metrics from these poses.

Per-video motion magnitude (computed independently for each video):

- *translation path length*: $\sum_{t=0}^{N-2} \|\boldsymbol{p}_{t+1} - \boldsymbol{p}_t\|$, the total distance the camera travels.
- *max translation*: $\max_t \|\boldsymbol{p}_t - \boldsymbol{p}_0\|$, the farthest the camera reaches from its starting position.
- *rotation total* (degrees): $\sum_{t=0}^{N-2} \angle\big(R_{t+1}R_t^\top\big)$, the accumulated frame-to-frame rotation magnitude.
- *max rotation* (degrees): $\max_t \angle\big(R_t R_0^\top\big)$, the largest deviation from the initial orientation.

Cross-seed trajectory diversity (computed per prompt across the $L = 8$ seeds, then averaged over prompts). Let $\boldsymbol{e}^{(i)} = \boldsymbol{p}_{N-1}^{(i)}$ be the endpoint of the i-th seed and $\{\boldsymbol{p}_t^{(i)}\}$ its full trajectory:

- *endpoint spread*: $\big\|\mathrm{std}_i \boldsymbol{e}^{(i)}\big\|_2$, the ℓ_2 norm of the per-axis standard deviation of the endpoints.
- *mean pairwise endpoint distance*: $\frac{2}{L(L-1)} \sum_{i<j} \|\boldsymbol{e}^{(i)} - \boldsymbol{e}^{(j)}\|$, averaged over all $\binom{L}{2} = 28$ seed pairs.
- *mean pairwise trajectory distance*: $\frac{2}{L(L-1)} \sum_{i<j} \frac{1}{N} \sum_t \|\boldsymbol{p}_t^{(i)} - \boldsymbol{p}_t^{(j)}\|$, the average per-frame distance between trajectories, capturing path-level (not just endpoint) differences.

E.3.4 VISIONREWARD

We follow the same VisionReward settings as Causal Forcing Zhu et al. (2026). We use VisionReward to evaluate overall visual quality. Each sub-score can be positive or negative and lies in $[-1, 1]$, where -1 indicates the worst quality and 1 the best. The final VisionReward score is computed as a weighted sum using the official weights. We additionally use VisionReward's prompt-alignment sub-score to evaluate instruction following by querying VisionReward with the official prompt: *Does the video meet some of the requirements stated in the text "[[prompt]]"?*

F ADDITIONAL RESULTS

F.1 ADDITIONAL RESULTS FOR TEXT-TO-VIDEO GENERATION

We show more visualization results for the text-to-video generation. And **We refer the readers to the supplementary material for full video comparisons**

F.1.1 DIVERSITY VISUALIZATION RESULTS

We further illustrate the diversity advantage of our method by generating videos with 8 different random seeds under the same text prompt. Fig. 7 shows the middle frame of each generated video for DMD2, DP-DMD, and ours. Each row corresponds to a different pair of seeds (rows 1–4 cover seeds 1–8). Within each method, the two columns show two different seeds; within each row,

the same two seeds are shown across the three methods. Compared to DMD2 and DP-DMD, our method produces clearly more diverse outputs across seeds, with larger variation in subject pose, layout, and scene composition. And **We refer the readers to the supplementary material for full video comparisons**



Figure 7: **Diversity visualization across 8 random seeds for the same prompt**. Columns 1–2: DMD2; columns 3–4: DP-DMD; columns 5–6: Ours. Each row shows the middle frame of two videos generated with different seeds. Our method produces visibly more diverse outputs across seeds.

F.2 Main Results for Autoregressive Video Generation

We refer the readers to the supplementary material for full video comparisons

F.3 Additional Results for Image-to-Video Generation

We refer the readers to the supplementary material for full video comparisons



Figure 8: **I2V results on the ViPE test set.** The colored outline indicates the input image. Our method produces videos that closely follow the first frame and remain coherent across the full sequence, whereas DMD2 and DP-DMD introduce structural anomalies such as an extra cowboy appearing from nowhere (highlighted by white circles).

F.4 LIMITATION

Limitations While DFD substantially improves few-step video generation, its performance remains limited under an extremely constrained generation budget. As shown in Fig. 9, when the generation is restricted to two steps or fewer, the two step distilled Cosmos-Predict2.5-2B model still struggles to produce high-quality videos: fast-moving objects, such as the man's hands, appear blurry; fine details, such as the woman's facial features, lose fidelity; and the model sometimes collapses to nearly static videos with little motion. These results suggest that, despite the improvements brought by DFD, generating high-quality and temporally dynamic videos remains challenging in the highly aggressive two-step-or-fewer regime.



Figure 9: **Limitation of DFD under a two-step generation budget.** With two-step distillation of the Cosmos-Predict2.5-2B model, DFD still produces blurry results for fast-moving content (e.g., the man's hands), loses fine detail (e.g., the woman's face), or collapses to highly static videos.

Text Prompts. We show part of the text prompts here:

Example Prompts

- Pikachu and Eevee dancing joyfully together on a dirt path in a bright, sunlit forest.
- Cheerful yellow cartoon dragon with green spikes and colorful wings, standing on a bright grassy hill.
- SpongeBob SquarePants dancing energetically inside his pineapple house, surrounded by floating bubbles.
- Cartoon-style family of polar bears playing on a floating ice floe.
- Animated squirrel gathering acorns in a colorful autumn forest.
- Anime-style mischievous cat batting at a ball of yarn in a cozy living room.
- Animation of a graceful fox running through a snowy landscape.
- Cartoon-style happy dog playing fetch on a bright sandy beach.
- Animated traditional Japanese village at sunset with glowing paper lanterns and falling cherry blossoms.
- Animated misty village with green-roofed wooden houses, a cobblestone path, and vibrant orange flowers.
- An animated princess and men in tuxedos interacting cheerfully by a grand clock tower at night.
- Cartoon-style elaborate underwater palace made of shells and coral.
- Animated complex European castle perched on a rugged, mist-shrouded mountain peak.
- Anime-style bustling Asian street market with colorful awnings, glowing signs, and dense pedestrian traffic.
- Cartoon-style quirky, multi-level treehouse with suspension bridges and tire swings in a forest.
- Animation of a historic lighthouse on a rocky coast battling powerful stormy waves.
- Anime-style Sailor Moon walking down a neon city street, looking back over her shoulder with a smile.
- Cheerful snowman standing on a snowy roof, looking down at a lit Christmas tree in a town square.

- Anime close-up of an angry girl with pink twin-tails, golden horns, and a dark purple outfit.
- Anime scene of a grinning Naruto holding ramen, standing next to a smiling girl with lavender hair.
- Cartoon blonde mermaid with a white tiara, swimming happily through a vibrant coral reef.
- Three animated women excitedly examining a map at a table in an ornate room.
- A cartoon woman smiling and sipping coffee from a red cup at a sunny, cozy café.
- A romantic animated couple gazing at each other under a Van Gogh-style swirling starry night sky.
- Two animated princesses having a magical outdoor tea party among lush flowers and glowing sparkles.
- Two animated princesses dancing joyfully on a palace terrace in the rain under a vibrant rainbow.
- An animated princess pointing enthusiastically toward a seated queen at a magical royal banquet.
- An animated princess pointing enthusiastically toward a seated queen at a magical royal banquet.
- Animation of a street artist creating a detailed vibrant mural on a brick wall.
- An animated woman in a blue gown gazing from a rose-adorned balcony overlooking a fairy-tale city.
- A cheerful cartoon princess gesturing joyfully in a vibrant, sunny garden with a distant palace.
- A tearful cartoon princess in a starry gown looking up at the night sky, with a majestic castle behind her.
- Cartoon-style towering waterfall plunging into a clear turquoise pool surrounded by jungle foliage.
- Animation of a vast desert with towering sand dunes stretching to the horizon under a blazing sun.
- Anime-style dense, ancient forest where beams of sunlight filter dramatically through the high canopy.
- Cartoon-style bustling cityscape viewed from a high rooftop, with lights blinking on as twilight descends.
- Animation of a tranquil lake perfectly reflecting the purple and orange colors of a vibrant dawn.
- Cartoon pig chef wearing a tall white hat, proudly holding a strawberry cake in a modern kitchen.
- Cartoon-style towering stack of pancakes dripping with maple syrup and melting butter in a cozy diner.
- Animation of a detailed chef's market stall with rows of colorful fresh fruits and vegetables.
- Anime-style detailed close-up of a steaming hot bento box full of intricate, colorful food.
-
- Cartoon-style conveyor belt with tiny, perfectly formed sushi plates moving past happy diners.
- Animated bakery window display filled with intricate pastries, artisanal bread, and glowing warmth.

- Cartoon-style colorful ice cream truck with children lining up on a hot summer day.
- Anime-style traditional tea ceremony performed in a peaceful garden with gentle steam rising from the cups.
- Cartoon-style bright red vintage steam train chugging across rolling green hills under a bright sky.
- Animated futuristic spaceship smoothly landing on a alien planet with purple vegetation.
- Whimsical stop-motion style yellow submarine exploring deep, glowing underwater ruins.
- Anime depiction of a sleek sports car speeding through a tunnel illuminated by passing neon lights.
- Cartoon-style biplane flying dynamic loop-the-loops amidst fluffy, stylized white clouds.
- Animated pirate ship with detailed sails navigating a vast, sun-drenched ocean at sunset.
- Cartoon delivery scooter expertly navigating a bustling, rainy city street.
- Anime-style flying motorcycle soaring high above a dense, green forest canopy.
- Cartoon-style dusty antique clock shop with dozens of clocks ticking and chiming simultaneously.
- Animated messy artist's desk covered in used paintbrushes, tubes of paint, and open sketchbooks.
- Whimsical stop-motion style pair of worn leather boots resting by a crackling fireplace.
-
- Anime-style close-up of a well-traveled backpack adorned with colorful keychains and buttons.
- Cartoon-style bookshelf packed with a vast, eclectic collection of books in a library corner.
- Animation of a grand grandfather clock with its large pendulum slowly swinging back and forth.
- Cartoon-style gardening tools—a trowel, gloves, and seed packets—resting on a rustic wooden bench.
- Anime-style glowing crystal globe positioned on a wooden desk within a wizard's library.
- Animated messy child's playroom filled with various toys and building blocks scattered on the floor.
- Anime-style depiction of a massive magical library where books fly autonomously between high shelves.
- Animation of a whimsical city made of clouds floating in the sky, connected by glowing bridges of light.
- Cartoon-style encounter with a friendly forest spirit in a glowing, ancient grove of mossy trees.
- Anime-style powerful wizard casting a complex spell that generates swirling vortexes of magical energy.
- Animation of a mischievous imp hiding behind colorful, glowing mushrooms within a deep cave.
- Cartoon-style friendly dragon and a knight playing chess peacefully under the shade of a large tree.

- Anime-style phoenix rising dramatically from a bed of glowing ashes with brilliant, fiery feathers.
- Animation of an intricate city built entirely from colorful crystals that hum and pulse with light and energy.
- Whimsical cartoon depiction of a celestial garden where stars grow like flowers on glowing vines.
- Whimsical cartoon depiction of a celestial garden where stars grow like flowers on glowing vines.
- A young woman with curly red hair laughs while spinning a yellow umbrella in slow motion as rain falls around her on a cobblestone street.
- Close-up of an elderly man's weathered hands carefully tying a fly fishing lure beside a misty mountain stream at dawn.
- Cinematic shot of a lone figure in a long coat walking across an endless salt flat under a vast purple twilight sky.
- A barista pulls a perfect espresso shot, dark crema swirling into a white ceramic cup in soft morning light.
- Drone footage soaring above a turquoise coastline where waves crash against jagged white cliffs.
- A chef in a black apron flips a sizzling steak in a cast iron pan, flames leaping briefly upward.
- Close-up of a honeybee landing on a sunflower, pollen visibly clinging to its legs as it crawls across the petals.
- A vintage motorcycle rolls slowly down an empty desert highway at sunset, heat shimmering off the asphalt.
- Cinematic tracking shot of a ballerina in a white tutu rehearsing alone in a sunlit wooden studio.
- Time-lapse of storm clouds rolling rapidly across a wheat field as the wind bends the golden stalks.
- A black cat slinks silently across a wet rooftop at night, eyes glowing under a streetlamp.
- Macro shot of a single dewdrop sliding down a green blade of grass at sunrise.
- A woman in a flowing red dress runs barefoot along a deserted beach, footprints quickly washed away by the surf.
- First-person view from inside a glass elevator rising rapidly up a futuristic skyscraper at night.
- A young boy lights sparklers in his backyard as twilight settles, his face illuminated in flickering gold.
- Aerial shot of a small wooden boat carving a wake across a glassy lake surrounded by dense pine forest.
- A potter's hands shape wet clay on a spinning wheel, water and clay slipping between their fingers.
- Cinematic close-up of an astronaut's helmet visor reflecting the curve of Earth from low orbit.
- A street musician plays a worn acoustic guitar on a busy sidewalk as commuters hurry past in motion blur.
- Slow-motion shot of a basketball swishing through a chain net on an outdoor court at golden hour.
- A fox steps cautiously into a snowy clearing at dusk, ears twitching as snowflakes drift down around it.

- Wide shot of a Tibetan monastery perched on a cliff edge, prayer flags whipping in the high mountain wind.
- A barista pours latte art into a cappuccino, viewed from directly above as a delicate rosetta forms.
- A child runs through a sprinkler in a suburban backyard on a hot summer afternoon, laughing.
- Macro footage of ink spreading and blooming through clear water, forming twisting black tendrils.
- Cinematic shot of a saxophonist playing on a foggy New Orleans street corner, neon signs glowing behind him.
- An elderly couple slow dances in their kitchen as warm afternoon light streams through gauzy curtains.
- Drone shot pulling back from a single red farmhouse to reveal endless rolling vineyards stretching to the horizon.
- A blacksmith hammers a glowing orange blade on an anvil, sparks flying with each strike.
- Underwater shot of a sea turtle gliding gracefully through shafts of sunlight piercing a coral reef.
- A businesswoman in a tailored suit strides confidently across a rain-soaked Tokyo crosswalk at night, neon reflecting in puddles.
- Close-up of a vinyl record dropping onto a turntable, the needle settling as it begins to spin.
- Time-lapse of a city skyline transitioning from sunset to night, windows lighting up across the buildings.
- A surfer paddles into a massive wave at dawn, silhouetted against a pink and orange sky.
- Cinematic shot of a horse galloping across a misty Scottish moor, hooves throwing up clumps of damp earth.
- A jazz pianist's hands move expressively across the keys in a dimly lit speakeasy filled with soft amber light.
- A baker dusts flour across a freshly braided loaf of bread on a wooden countertop in a sunlit kitchen.
- Slow-motion shot of a tennis ball striking a racket, the strings flexing dramatically on impact.
- A father teaches his young daughter to ride a bicycle along a tree-lined neighborhood street in autumn.
- Aerial view of a fishing village at dawn, colorful boats returning to harbor across calm waters.
- Macro shot of a butterfly slowly opening and closing its wings on a lavender stem in a summer breeze.
- A man reads a leather-bound book in a worn armchair beside a crackling fireplace as snow falls past the window.
- Cinematic shot of a train emerging from a tunnel into a sunlit mountain valley filled with wildflowers.
- A surgeon in scrubs washes their hands at a hospital sink, water cascading over their fingers in clinical white light.
- Close-up of raindrops striking a car windshield, the wipers sweeping rhythmically across the glass.

- A florist arranges a bouquet of peonies and eucalyptus on a marble counter in a bright minimalist shop.
- A lone hiker reaches a mountain summit at sunrise and stands silhouetted against the rising sun, arms outstretched.
- Slow-motion shot of a glass shattering as it strikes a tile floor, shards spreading outward.
- Drone footage following a red kayak winding through a narrow emerald-green river canyon.
- A young man in a leather jacket leans against a vintage car at a 1950s-style diner, neon signs buzzing overhead.
- Cinematic shot of two dancers performing a tango in an empty warehouse lit by a single overhead bulb.
- A child's hand reaches up to touch the glass of an aquarium as a stingray glides past on the other side.
- Wide shot of a hot air balloon festival at dawn, dozens of colorful balloons inflating across a misty field.
- A glassblower shapes molten glass at the end of a long iron rod, the orange glow lighting their focused face.
- Close-up of espresso dripping into a clear glass mug, steam curling upward in soft window light.
- A street artist spray-paints a vibrant mural on a brick wall in an alley, bystanders watching in the background.
- Slow-motion shot of a diver leaping from a high cliff into a deep blue ocean below.
- A woman pours steaming tea from a cast iron kettle into a delicate porcelain cup at a wooden table.
- Aerial shot of an autumn forest in full color, a winding road cutting through a fiery red and orange canopy.
- A construction worker welds a steel beam high above a city skyline, sparks raining down against the sunset.
- Cinematic shot of a vintage typewriter on a desk by a rainy window, paper slowly being pulled from the carriage.
- A skateboarder grinds along a metal handrail in an urban plaza, captured in slow motion at golden hour.
- A grandmother teaches her granddaughter to roll dumplings at a flour-dusted kitchen counter, both smiling.
- Underwater shot of a school of silver fish parting around a diver swimming through a sunlit kelp forest.
- Time-lapse of a flower blooming from a tight bud to a full open blossom over the course of a day.
- A scientist examines a glowing blue liquid in a test tube inside a softly lit laboratory.
- Cinematic wide shot of a desert caravan crossing massive sand dunes at sunset, long shadows trailing behind.
- A coffee shop window steams up from the inside as rain streaks down the outside, a person reading at the counter.
- Slow-motion shot of a flock of birds taking off from a marsh at dawn, wings beating against pink mist.
- A young woman pins a Polaroid photo onto a string of lights in her bedroom, the photo gently swinging.
- Drone footage gliding low over a frozen lake at dusk, a single ice skater carving graceful arcs across its surface.

F.5 Societal Impacts

Our work on distilling video diffusion models presents both positive and negative potential societal impacts. On the positive side, our method significantly reduces the computational cost and inference time required for high-fidelity video generation. This democratizes access to creative tools, allowing users with limited compute resources to generate videos, while also reducing the carbon footprint and energy consumption associated with deploying large-scale generative models. On the negative side, accelerating video generation inherently scales the risks associated with the base models. Faster generation can facilitate the rapid creation of deepfakes, misinformation, and malicious content. While our theoretical framework improves algorithmic efficiency, it does not inherently prevent these misuses. Safe deployment of these distilled models in real-world applications will require coupling them with robust safety filters, provenance tracking, and watermarking techniques.