
T2V-Turbo: Breaking the Quality Bottleneck of Video Consistency Model with Mixed Reward Feedback

Jiachen Li¹, Weixi Feng¹, Tsu-Jui Fu¹, Xinyi Wang¹, Sugato Basu²,
Wenhu Chen³, William Yang Wang¹

¹UC Santa Barbara, ²Google, ³University of Waterloo

¹{jiachen_li, weixifeng, tsu-juifu, xinyi_wang, william}@cs.ucsb.edu
²sugato@google.com ³wenhuchen@uwaterloo.ca

Project Page: <https://t2v-turbo.github.io>

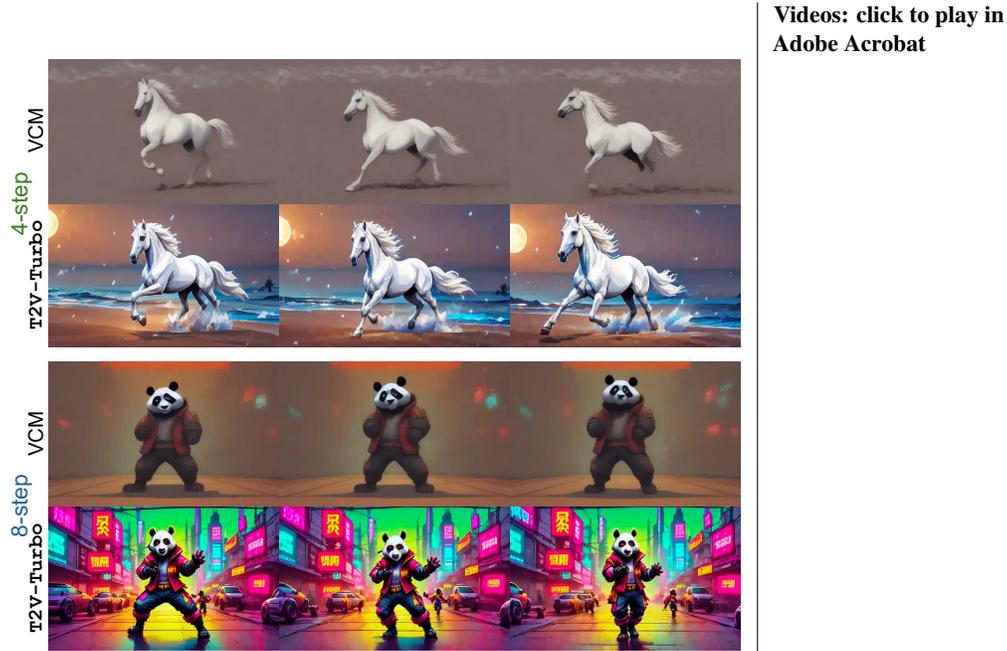


Figure 1: By integrating reward feedback during consistency distillation from VideoCrafter2 [Chen et al., 2024], our T2V-Turbo (VC2) can generate high-quality videos with 4-8 inference steps, breaking the quality bottleneck of a VCM [Wang et al., 2023a]. Appendix F includes the corresponding text prompts.

Abstract

Diffusion-based text-to-video (T2V) models have achieved significant success but continue to be hampered by the slow sampling speed of their iterative sampling processes. To address the challenge, consistency models have been proposed to facilitate fast inference, albeit at the cost of sample quality. In this work, we aim to break the quality bottleneck of a video consistency model (VCM) to achieve **both fast and high-quality video generation**. We introduce T2V-Turbo, which integrates feedback from a mixture of differentiable reward models into the consistency distillation (CD) process of a pre-trained T2V model. Notably, we directly optimize rewards associated with single-step generations that arise naturally from computing the CD loss, effectively bypassing the memory constraints imposed by backpropagating gradients through an iterative sampling process. Remarkably, the 4-step generations from our T2V-Turbo achieve the highest total score on VBench [Huang et al., 2024], even surpassing Gen-2 [Esser et al., 2023] and

Pika [Pika Labs, 2023]. We further conduct human evaluations to corroborate the results, validating that the 4-step generations from our T2V-Turbo are preferred over the 50-step DDIM samples from their teacher models, representing more than a tenfold acceleration while improving video generation quality.

1 Introduction

Diffusion model (DM) [Sohl-Dickstein et al., 2015, Ho et al., 2020] has emerged as a powerful framework for neural image [Betker et al., 2023, Rombach et al., 2022, Esser et al., 2024, Saharia et al., 2022] and video synthesis [Singer et al., 2022, Ho et al., 2022a, He et al., 2022, Wang et al., 2023b, Zhang et al., 2023], leading to the development of cutting-edge text-to-video (T2V) models like Sora [Brooks et al., 2024], Gen-2 [Esser et al., 2023] and Pika [Pika Labs, 2023]. Although the iterative sampling process of these diffusion-based models ensures high-quality generation, it significantly slows down inference, hindering their real-time applications. On the other hand, existing open-sourced T2V models including VideoCrafter [Chen et al., 2023, 2024] and ModelScopeT2V [Wang et al., 2023c] are trained on web-scale video datasets, e.g., WebVid-10M [Bain et al., 2021], with varying video qualities. Consequently, the generated videos often appear visually unappealing and fail to align accurately with the text prompts, deviating from human preferences.

Efforts have been made to address the issues listed above. To accelerate the inference process, Wang et al. [2023a] applies the theory of consistency distillation (CD) [Song et al., 2023, Song and Dhariwal, 2023, Luo et al., 2023a] to distill a video consistency model (VCM) from a teacher T2V model, enabling plausible video generations in just 4-8 inference steps. However, the quality of VCM’s generations is naturally bottlenecked by the performance of the teacher model, and the reduced number of inference steps further diminishes its generation quality. On the other hand, to align generated videos with human preferences, InstructVideo [Yuan et al., 2023] draws inspiration from image generation techniques [Dong et al., 2023, Clark et al., 2023, Prabhudesai et al., 2023] and proposes backpropagating the gradients of a differentiable reward model (RM) through the iterative video sampling process. However, calculating the full reward gradient is prohibitively expensive, resulting in substantial memory costs. Consequently, InstructVideo truncates the sampling chain by limiting gradient calculation to only the final DDIM step, compromising optimization accuracy. Additionally, InstructVideo is limited by its reliance on an image-text RM, which fails to fully capture the transition dynamic of a video. Empirically, InstructVideo only conducts experiments on a limited set of user prompts, the majority of which are related to animals. As a result, its generalizability to a broader range of prompts remains unknown.

In this paper, we aim to achieve fast and high-quality video generation by breaking the quality bottleneck of a VCM. We introduce T2V-Turbo, which integrates reward feedback from a mixture of RMs into the process of distilling a VCM from a teacher T2V model. Besides utilizing an image-text RM to align individual video frames with human preference, we further incorporate reward feedback

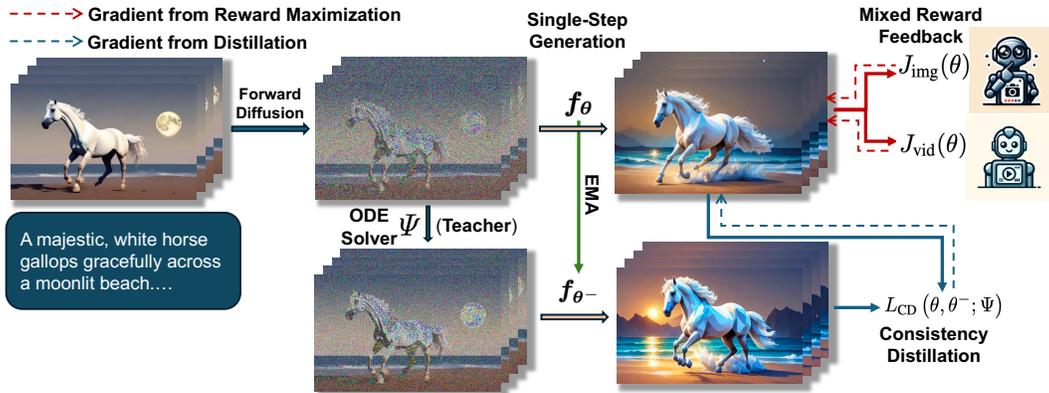


Figure 2: Overview of the training pipeline of our T2V-Turbo. We integrate reward feedback from both an image-text RM and a video-text RM into the VCD procedures by backpropagating gradient through the single-step generation process of our T2V-Turbo.

from a video-text RM to comprehensively evaluate the temporal dynamics and transitions in the generated videos. We highlight that our reward optimization avoids tackling the highly memory-intensive issues associated with backpropagating gradients through an iterative sampling process. Instead, we directly optimize rewards of the single-step generations that arise from computing the CD loss, effectively bypassing the memory constraints faced by conventional methods that optimize a DM [Yuan et al., 2023, Xu et al., 2024a, Clark et al., 2023, Prabhudesai et al., 2023].

Empirically, we demonstrate the superiority of our T2V-Turbo in generating high-quality videos within 4-8 inference steps. To illustrate the applicability of our methods, we distill T2V-Turbo (VC2) and T2V-Turbo (MS) from VideoCrafter2 [Chen et al., 2024] and ModelScopeT2V [Wang et al., 2023c], respectively. Remarkably, the 4-step generation results from both variants of our T2V-Turbo outperform SOTA models on the video evaluation benchmark VBench [Huang et al., 2024], even surpassing proprietary systems such as Gen-2 [Esser et al., 2023] and Pika [Pika Labs, 2023] that are trained with extensive resources. We further corroborate the results by conducting human evaluation using 700 prompts from the EvalCrafter [Liu et al., 2023] benchmark, validating that the 4-step generations from T2V-Turbo are favored by human over the 50-step DDIM samples from their teacher T2V models, which represents over tenfold inference acceleration and enhanced video generation quality.

Our contributions are threefold:

- Learn a T2V model with feedback from a mixture of RMs, including a video-text model. To the best of our knowledge, we are the first to do so.
- Establish a new SOTA on the VBench with only 4 inference steps, outperforming proprietary models trained with substantial resources.
- 4-step generations from our T2V-Turbo are favored over the 50-step generation from its teacher T2V model as evidenced by human evaluation, representing over 10 times inference acceleration with quality improvement.

2 Preliminaries

Diffusion models (DMs). In the forward process, DMs progressively inject Gaussian noise into the original data distribution $p_{\text{data}}(\mathbf{x}) \equiv p_0(\mathbf{x}_0)$ and perturb it into a marginal distribution $p_t(\mathbf{x}_t)$ with the transition kernel $p_{0t}(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t|\alpha(t)\mathbf{x}_0, \beta^2(t)\mathbf{I})$ at timestep t . $\alpha(t)$ and $\beta(t)$ correspond to the noise schedule. In the reverse process, DMs sequentially recover the data from a noise sampled from the prior distribution $p_T(\mathbf{x}_T) := \mathcal{N}(\mathbf{x}_T|\mathbf{0}, \beta^2(T)\mathbf{I})$. The reverse-time SDE can be modeled by an ordinary differential equation (ODE), known as the Probability Flow (PF-ODE) [Song et al., 2020a]:

$$d\mathbf{x}_t = \left[\boldsymbol{\mu}(t) \mathbf{x}_t - \frac{1}{2} \sigma(t)^2 \nabla \log p_t(\mathbf{x}_t) \right] dt, \quad \mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \beta^2(T)\mathbf{I}). \quad (1)$$

where $\boldsymbol{\mu}(\cdot)$ and $\sigma(\cdot)$ are the drift and diffusion coefficients, respectively, with the following properties:

$$\boldsymbol{\mu}(t) = \frac{d \log \alpha(t)}{dt}, \quad \sigma^2(t) = \frac{d\beta^2(t)}{dt} - 2 \frac{d \log \alpha(t)}{dt} \beta^2(t). \quad (2)$$

The PF-ODE’s solution trajectories, when sampled at any timestep t , align with the distribution $p_t(\mathbf{x}_t)$. Empirically, a denoising model $\epsilon_\theta(\mathbf{x}_t, t)$ is trained to approximate the score function $-\nabla \log p_t(\mathbf{x}_t)$ via score matching. During the sampling phase, one begins with a sample $\mathbf{x}_T \sim p_T(\mathbf{x}_T)$ and follows the empirical PF-ODE below to obtain a sample $\hat{\mathbf{x}}_0$.

$$d\mathbf{x}_t = \left[\boldsymbol{\mu}(t) \mathbf{x}_t + \frac{1}{2} \sigma(t)^2 \epsilon_\theta(\mathbf{x}_t, t) \right] dt, \quad \mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \beta^2(T)\mathbf{I}). \quad (3)$$

In this paper, we focus on diffusion-based T2V models, which operate on the video latent space \mathcal{Z} and train a denoising model $\epsilon_\theta(\mathbf{z}_t, \mathbf{c}, t)$ conditioned on the text prompt \mathbf{c} , where \mathbf{z}_t is obtained by perturbing the image latent $\mathbf{z} = \mathcal{E}(\mathbf{x})$, $\in \mathcal{Z}$ and \mathcal{E} is a VAE [Kingma and Welling, 2013] encoder. The T2V models employ Classifier-Free Guidance (CFG) [Ho and Salimans, 2021] to enhance the quality of conditional sampling by substituting the noise prediction with a linear combination of conditional and unconditional noise predictions for denoising, i.e., $\tilde{\epsilon}_\theta(\mathbf{z}_t, \omega, \mathbf{c}, t) = (1+\omega)\epsilon_\theta(\mathbf{z}_t, \mathbf{c}, t) - \omega\epsilon_\theta(\mathbf{z}_t, \emptyset, t)$, where ω is the CFG scale. After the completion of the inference process, we can generate a video by $\hat{\mathbf{x}}_0 = \mathcal{D}(\mathbf{z}_0)$ with the VAE decoder \mathcal{D} corresponding to \mathcal{E} .

Consistency Distillation. Conventional methods [Ho et al., 2020, Song et al., 2020b] generate their samples by solving the PF-ODE sequentially, leading to DM’s slow inference speed. To tackle this problem, *consistency models* (CM) [Song et al., 2023, Song and Dhariwal, 2023] propose to learn a consistency function $\mathbf{f} : (\mathbf{x}_t, t) \mapsto \mathbf{x}_\epsilon$ to directly map any \mathbf{x}_t on the PF-ODE trajectory to its origin, where ϵ is a fixed small positive number. And thus, the consistency function \mathbf{f} has the following *self-consistency* property

$$\mathbf{f}(\mathbf{x}_t, t) = \mathbf{f}(\mathbf{x}'_{t'}, t'), \forall t, t' \in [\epsilon, T], \quad (4)$$

where \mathbf{x}_t and $\mathbf{x}'_{t'}$ are from the same PF-ODE. We can model \mathbf{f} with a CM \mathbf{f}_θ . When tackling the PF-ODE of a T2V model that operates on the video latent space \mathcal{Z} , we aim to learn a video consistency model (VCM) [Luo et al., 2023a, Wang et al., 2023a] $\mathbf{f}_\theta : (\mathbf{z}_t, \omega, \mathbf{c}, t) \mapsto \mathbf{z}_0 \in \mathcal{Z}$. To ensure $\mathbf{f}_\theta(\mathbf{z}, \omega, \mathbf{c}, t) = \mathbf{z}$, we parameterize \mathbf{f}_θ as

$$\mathbf{f}_\theta(\mathbf{z}, \omega, \mathbf{c}, t) = c_{\text{skip}}(t)\mathbf{z} + c_{\text{out}}(t)F_\theta(\mathbf{z}, \omega, \mathbf{c}, t), \quad (5)$$

where $c_{\text{skip}}(t)$ and $c_{\text{out}}(t)$ are differentiable functions with $c_{\text{skip}}(\epsilon) = 1$ and $c_{\text{out}}(\epsilon) = 0$, and F_θ is modeled as a neural network. We can distill a \mathbf{f}_θ from a pre-trained T2V DM by minimizing the *consistency distillation* (CD) [Song et al., 2023, Luo et al., 2023a] loss as below

$$L_{\text{CD}}(\theta, \theta^-; \Psi) = \mathbb{E}_{\mathbf{z}, \mathbf{c}, \omega, n} \left[d \left(\mathbf{f}_\theta(\mathbf{z}_{t_{n+k}}, \omega, \mathbf{c}, t_{n+k}), \mathbf{f}_{\theta^-}(\hat{\mathbf{z}}_{t_n}^{\Psi, \omega}, \omega, \mathbf{c}, t_n) \right) \right], \quad (6)$$

where $d(\cdot, \cdot)$ is a distance function. θ^- is updated by the exponential moving average (EMA) of θ , i.e., $\theta^- \leftarrow \text{stop_grad}(\mu\theta + (1 - \mu)\theta^-)$. $\hat{\mathbf{z}}_{t_n}^{\Psi, \omega}$ is an estimate of \mathbf{z}_{t_n} obtained by the numerical augmented PF-ODE solver Ψ parameterized by ψ and k is the skipping interval

$$\hat{\mathbf{z}}_{t_n}^{\Psi, \omega} \leftarrow \mathbf{z}_{t_{n+k}} + (1 + \omega)\Psi(\mathbf{z}_{t_{n+k}}, t_{n+k}, t_n, \mathbf{c}; \psi) - \omega\Psi(\mathbf{z}_{t_{n+k}}, t_{n+k}, t_n, \emptyset; \psi). \quad (7)$$

We follow the LCM paper [Luo et al., 2023a] to use DDIM [Song et al., 2020b] as the ODE solver Ψ and defer the formula of the DDIM solver to Appendix A.

3 Training T2V-Turbo with Mixed Reward Feedback

In this section, we present the training pipeline to derive our T2V-Turbo. To facilitate fast and high-quality video generation, we integrate reward feedback from multiple RMs into the LCD process when distilling from a teacher T2V model. Figure 2 provides an overview of our framework. Notably, we directly leverage the single-step generation $\hat{\mathbf{z}}_0 = \mathbf{f}_\theta(\mathbf{z}_{t_{n+k}}, \omega, \mathbf{c}, t_{n+k})$ arise from computing the CD loss L_{CD} (6) and optimize the video $\hat{\mathbf{x}}_0 = \mathcal{D}(\hat{\mathbf{z}}_0)$ decoded from it towards multiple differentiable RMs. As a result, we avoid the challenges associated with backpropagating gradients through an iterative sampling process, which is often confronted by conventional methods optimizing DMs [Clark et al., 2023, Xu et al., 2024a, Yuan et al., 2023].

In particular, we leverage reward feedback from an image-text RM to improve human preference on each individual video frame (Sec. 3.1) and further utilize the feedback from a video-text RM to improve the temporal dynamics and transitions in the generated video (Sec. 3.2).

3.1 Optimizing Human Preference on Individual Video Frames

Chen et al. [2024] achieve high-quality video generation by including high-quality images as single-frame videos when training the T2V model. Inspired by their success, we align each individual video frame with human preference by optimizing towards a differentiable image-text RM \mathcal{R}_{img} . In particular, we randomly sample a batch of M frames $\{\hat{\mathbf{x}}_0^1, \dots, \hat{\mathbf{x}}_0^M\}$ from the decoded video $\hat{\mathbf{x}}_0$ and maximize their scores evaluated by \mathcal{R}_{img} as below

$$J_{\text{img}}(\theta) = \mathbb{E}_{\hat{\mathbf{x}}_0, \mathbf{c}} \left[\sum_{m=1}^M \mathcal{R}_{\text{img}}(\hat{\mathbf{x}}_0^m, \mathbf{c}) \right], \quad \hat{\mathbf{x}}_0 = \mathcal{D}(\mathbf{f}_\theta(\mathbf{z}_{t_{n+k}}, \omega, \mathbf{c}, t_{n+k})). \quad (8)$$

3.2 Optimizing Video-Text Feedback Model

Existing image-text RMs [Wu et al., 2023a, Xu et al., 2024a, Kirstain et al., 2024] are limited to assessing the alignment between individual video frames and the text prompt and thus cannot evaluate

through the temporal dimensions that involve inter-frame dependencies, such as motion dynamic and transitions [Huang et al., 2024, Liu et al., 2023]. To address these shortcomings, we further leverage a video-text RM \mathcal{R}_{vid} to assess the generated videos. The corresponding objective J_{vid} is given below

$$J_{\text{vid}}(\theta) = \mathbb{E}_{\hat{\mathbf{x}}_0, \mathbf{c}} [\mathcal{R}_{\text{vid}}(\hat{\mathbf{x}}_0, \mathbf{c})], \quad \hat{\mathbf{x}}_0 = \mathcal{D}(\mathbf{f}_{\theta}(\mathbf{z}_{t_{n+k}}, \omega, \mathbf{c}, t_{n+k})). \quad (9)$$

3.3 Summary

To this end, we can define the total learning loss L of our training pipeline as a linear combination of the L_{CD} in (6), J_{img} in (8), and J_{vid} in (9) with weighting parameters β_{img} and β_{vid} .

$$L(\theta, \theta^-; \Psi) = L_{\text{CD}}(\theta, \theta^-; \Psi) - \beta_{\text{img}} J_{\text{img}}(\theta) - \beta_{\text{vid}} J_{\text{vid}}(\theta) \quad (10)$$

To reduce memory and computational cost, we initialize our T2V-Turbo with the teacher model and only optimize the LoRA weights [Hu et al., 2021, Luo et al., 2023b] instead of performing full model training. After completing the training, we merge the LoRA weights so that the per-step inference cost of our T2V-Turbo remains identical to the teacher model. We include pseudo-codes for our training algorithm in Appendix B.

4 Experimental Results

Our experiments aim to demonstrate our T2V-Turbo’s ability to generate high-quality videos with 4-8 inference steps. We first conduct automatic evaluations on the standard benchmark VBench [Huang et al., 2024] to comprehensively evaluate our methods from various dimensions (Sec. 4.1) against a broad array of baseline methods. We then perform human evaluations with 700 prompts from the EvalCrafter [Liu et al., 2023] to compare the 4-step and 8-step generations from our T2V-Turbo with the 50-step generations from the teacher T2V models as well as the 4-step generations from the baseline VCM (Sec. 4.2). Finally, we perform ablation studies on critical design choices (Sec. 4.3).

Settings. We train T2V-Turbo (VC2) and T2V-Turbo (MS) by distilling from the teacher diffusion-based T2V models VideoCrafter2 [Chen et al., 2024] and ModelScopeT2V [Wang et al., 2023c], respectively. Similar to both teacher models, we conduct our training using the WebVid10M [Bain et al., 2021] datasets. We train our models on 8 NVIDIA A100 GPUs for 10K gradient steps without gradient accumulation. We set the batch size of training videos to 1 for each GPU device. We employ HPSv2.1 [Wu et al., 2023a] as our image-text RM \mathcal{R}_{img} . When distilling from VideoCrafter2, we utilize the 2nd Stage model of InternVideo2 (InternVid2 S2) [Wang et al., 2024] as our video-text RM \mathcal{R}_{vid} . When distilling from ModelScopeT2V, we set \mathcal{R}_{vid} to be ViCLIP [Wang et al., 2023d]. To optimize J_{img} (8), we randomly sample 6 frames from the video by setting $M = 6$. For the hyperparameters (HP), we set learning rate $1e - 5$ and guidance scale range $[\omega_{\min}, \omega_{\max}] = [5, 15]$. We use DDIM [Song et al., 2020b] as our ODE solver Ψ and set the skipping step $k = 20$. For T2V-Turbo (VC2), we set $\beta_{\text{img}} = 1$ and $\beta_{\text{vid}} = 2$. For T2V-Turbo (MS), we set $\beta_{\text{img}} = 2$ and $\beta_{\text{vid}} = 3$. We include further training details in Appendix A.

4.1 Automatic Evaluation on VBench

We evaluate our T2V-Turbo (VC2) and T2V-Turbo (MS) on the standard video evaluation benchmark VBench [Huang et al., 2024] to compare against a wide array of baseline methods. VBench is designed to comprehensively evaluate T2V models from 16 disentangled dimensions. Each dimension in VBench is tailored with specific prompts and evaluation methods.

Table 1 compares the 4-step generation of our methods with various baselines from the VBench leaderboard¹, including Gen-2 [Esser et al., 2023], Pika [Pika Labs, 2023], VideoCrafter1 [Chen et al., 2023], VideoCrafter2 [Chen et al., 2024], Show-1 [Zhang et al., 2023], LaVie [Wang et al., 2023b], and ModelScopeT2V [Wang et al., 2023c]. Table 4 in Appendix further compares our methods with VideoCrafter0.9 [He et al., 2022], LaVie-Interpolation [Wang et al., 2023b], Open-Sora [Open-Sora, 2024], and CogVideo [Hong et al., 2022]. The performance of each baseline method is directly reported from the VBench leaderboard. To obtain the results of our methods, we carefully follow VBench’s evaluation protocols by generating 5 videos for each prompt to calculate the metrics. We

¹https://huggingface.co/spaces/Vchitect/VBench_Leaderboard

Table 1: **Automatic Evaluation on VBench** [Huang et al., 2024]. We compare our T2V-Turbo (VC2) and T2V-Turbo (MS) with baseline methods across the 16 VBench dimensions. A higher score indicates better performance for a particular dimension. We bold the best results for each dimension and underline the second-best result. **Quality Score** is calculated with the 7 dimensions from the top table. **Semantic Score** is calculated with the 9 dimensions from the bottom table. **Total Score** a weighted sum of **Quality Score** and **Semantic Score**. Further details can be found in Appendix C. Both our T2V-Turbo (VC2) and T2V-Turbo (MS) **surpass all baseline methods with 4 inference steps** in terms of Total Score, including the proprietary systems Gen-2 and Pika.

Models	Total Score	Quality Score	Subject Consist.	BG Consist.	Temporal Flicker.	Motion Smooth.	Aesthetic Quality	Dynamic Degree	Image Quality
ModelScopeT2V	75.75	78.05	89.87	95.29	98.28	95.79	52.06	66.39	58.57
LaVie	77.08	78.78	91.41	97.47	98.30	96.38	54.94	49.72	61.90
Show-1	78.93	80.42	95.53	98.02	99.12	98.24	57.35	44.44	58.66
VideoCrafter1	79.72	81.59	95.10	98.04	98.93	95.67	62.67	55.00	65.46
Pika	80.40	82.68	96.76	98.95	99.77	<u>99.51</u>	<u>63.15</u>	37.22	62.33
VideoCrafter2	80.44	82.20	<u>96.85</u>	98.22	98.41	97.73	63.13	42.50	67.22
Gen-2	80.58	82.47	97.61	97.61	<u>99.56</u>	99.58	66.96	18.89	67.42
VCM (MS)	75.84	78.80	93.06	97.30	98.51	98.00	48.99	46.11	61.98
Our T2V-Turbo (MS)	<u>80.62</u>	82.15	94.82	<u>98.71</u>	97.99	95.64	60.04	66.39	<u>68.09</u>
VCM (VC2)	73.97	78.54	94.02	96.05	99.06	98.84	54.56	42.50	52.72
Our T2V-Turbo (VC2)	81.01	<u>82.57</u>	96.28	97.02	97.48	97.34	63.04	49.17	72.49

Models	Semantic Score	Object Class	Multiple Objects	Human Action	Color	Spatial Relation.	Scene	Appear. Style	Temporal Style	Overall Consist.
ModelScopeT2V	66.54	82.25	38.98	92.40	81.72	33.68	39.26	23.39	25.37	25.67
LaVie	70.31	91.82	33.32	96.80	86.39	34.09	52.69	23.56	<u>25.93</u>	26.41
Show-1	72.98	93.07	45.47	95.60	86.35	53.50	47.03	23.06	25.28	27.46
VideoCrafter1	72.22	78.18	45.66	91.60	93.32	58.86	43.75	24.41	25.54	26.76
Pika	71.26	87.45	46.69	88.00	85.31	<u>65.65</u>	44.80	21.89	24.44	25.47
VideoCrafter2	73.42	92.55	40.66	95.00	<u>92.92</u>	35.86	<u>55.29</u>	<u>25.13</u>	25.84	28.23
Gen-2	73.03	90.92	<u>55.47</u>	89.20	89.49	66.91	48.91	19.34	24.12	26.17
VCM (MS)	63.98	83.18	24.85	87.20	85.72	31.57	42.44	23.20	23.30	24.18
Our T2V-Turbo (MS)	<u>74.47</u>	<u>93.34</u>	58.63	95.80	89.67	45.74	48.47	23.23	25.92	27.51
VCM (VC2)	55.66	63.97	10.81	82.60	79.12	23.06	18.49	25.29	22.31	25.15
Our T2V-Turbo (VC2)	74.76	93.96	54.65	95.20	89.90	38.67	55.58	24.42	25.51	<u>28.16</u>

further train VCM (VC2) and VCM (MS) by distilling from VideoCrafter2 and ModelScopeT2V, respectively, without incorporating reward feedback, and then compare their results.

VBench has developed its own rules to calculate the **Total Score**, **Quality Score**, and **Semantic Score**. **Quality Score** is calculated with the 7 dimensions from the top table. **Semantic Score** is calculated with the 9 dimensions from the bottom table. And **Total Score** is a weighted sum of Quality Score and Semantic Score. Appendix C provides further details, including explanations for each dimension of VBench. As shown in Table 1, the 4-step generations of both our T2V-Turbo (MS) and T2V-Turbo (VC2) surpass all baseline methods on VBench in terms of Total Score. These results are particularly remarkable given that we even outperform the proprietary systems Gen-2 and Pika, which are trained with extensive resources. Even when distilling from a less advanced teacher model, ModelScopeT2V, our T2V-Turbo (MS) attains the second-highest Total Score, just below our T2V-Turbo (VC2). Additionally, our T2V-Turbo breaks the quality bottleneck of a VCM by outperforming its teacher T2V model, significantly improving over the baseline VCM.

4.2 Human Evaluation with 700 EvalCrafter Prompts

To verify the effectiveness of our T2V-Turbo, we compare the 4-step and 8-step generations from our T2V-Turbo with the 50-step DDIM samples from the corresponding teacher T2V models. We further compare the 4-step generations between our T2V-Turbo and their baseline VCMs when distilled from the same teacher T2V model. We leverage the 700 prompts from the EvalCrafter [Liu et al., 2023] video evaluation benchmark, which are constructed based on real-world user data.

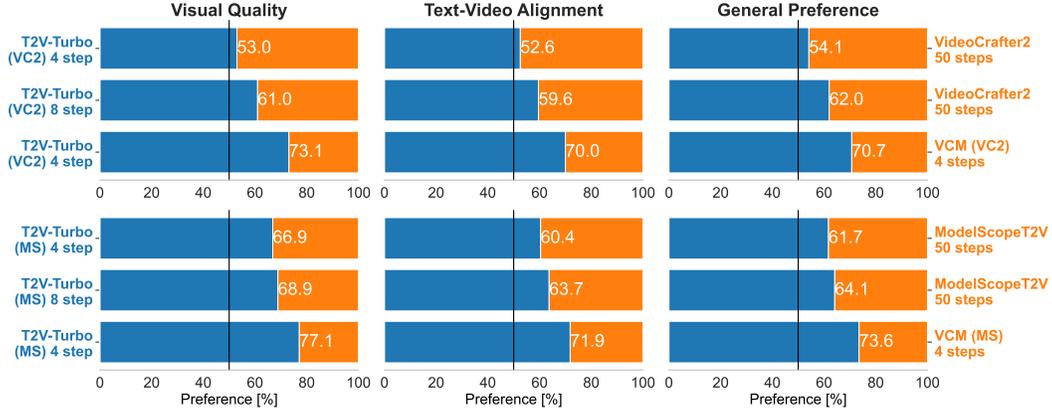


Figure 3: Human evaluation results with the 700 prompts from EvalCrafter [Liu et al., 2023]. We compare the 4-step and 8-step generations from our T2V-Turbo with their teacher T2V model and their baseline VCM. **Top**: results for T2V-Turbo (VC2). **Bottom**: results for T2V-Turbo (MS).

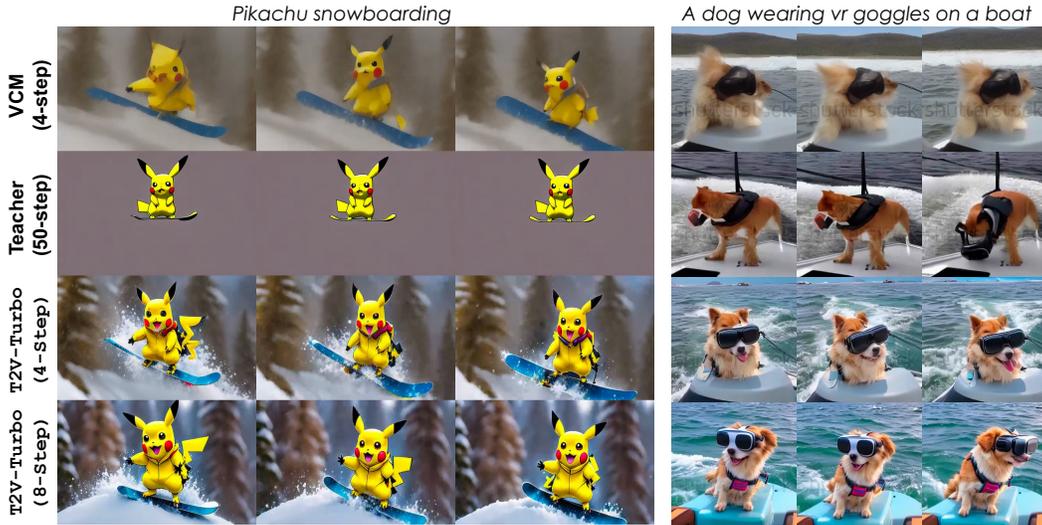


Figure 4: Qualitative comparisons between the 4-step VCM, 50-step teacher T2V, 4-step T2V-Turbo and 8-step T2V-Turbo generations. **Left**: (VC2), **Right**: (MS).

We hire human annotators from Amazon Mechanical Turk to compare videos generated from different models given the same prompt. For each comparison, the annotators need to answer three questions: Q1) Which video is more visually appealing? Q2) Which video better fits the text description? Q3) Which video do you prefer given the prompt? Appendix D includes additional details about how we set up the human evaluations.

Figure 3 provides the full human evaluation results. We also qualitatively compare different methods in Figure 4. Due to limited space, we include additional qualitative comparison results in Appendix F. Notably, the 4-step generations from our T2V-Turbo are favored by humans over the 50-step generation from their teacher T2V model, representing a 12.5 times inference acceleration with improving performance. By increasing the inference steps to 8, we can further improve the visual quality and text-video alignment of videos generated from our T2V-Turbo, reflected by the fact that our 8-step generations are more likely to be favored by the human compared to our 4-step generations in terms of all 3 evaluated metrics. Additionally, our T2V-Turbo significantly outperforms its baseline VCM, demonstrating the effectiveness of our methods, which incorporate a mixture of reward feedback into the model training.

Table 2: Ablation studies on the effectiveness of \mathcal{R}_{img} and \mathcal{R}_{vid} . We bold the highest score for each dimension for methods with the same teacher model. While incorporating feedback from \mathcal{R}_{img} is effective at improving both Quality Score and Semantic Score, integrating reward feedback from \mathcal{R}_{vid} can further improve the semantic score.

Models	Total Score	Quality Score	Subject Consist.	BG Consist.	Temporal Flicker.	Motion Smooth.	Aesthetic Quality	Dynamic Degree	Image Quality
VCM (MS)	75.84	78.80	93.06	97.30	98.51	98.00	48.99	46.11	61.98
VCM (MS) + \mathcal{R}_{vid}	77.28	78.76	93.24	97.67	98.49	97.27	51.70	55.00	56.40
VCM (MS) + \mathcal{R}_{img}	79.51	81.81	97.64	99.59	98.46	95.83	64.69	38.33	68.86
Our T2V-Turbo (MS)	80.62	82.15	94.82	98.71	97.99	95.64	60.04	66.39	68.09
VCM (VC2)	73.97	78.54	94.02	96.05	99.06	98.84	54.56	42.50	52.72
VCM (VC2) + \mathcal{R}_{vid}	77.57	80.08	95.46	96.69	98.78	98.79	58.66	25.00	65.75
VCM (VC2) + \mathcal{R}_{img}	80.42	82.59	96.52	97.31	97.50	97.29	63.08	47.50	72.91
Our T2V-Turbo (VC2)	81.01	82.57	96.28	97.02	97.48	97.34	63.04	49.17	72.49

Models	Semantic Score	Object Class	Multiple Objects	Human Action	Color	Spatial Relation.	Scene	Appear. Style	Temporal Style	Overall Consist.
VCM (MS)	63.98	83.18	24.85	87.20	85.72	31.57	42.44	23.20	23.30	24.18
VCM (MS) + \mathcal{R}_{vid}	71.35	91.14	45.64	94.60	86.97	39.74	48.55	22.90	25.91	26.81
VCM (MS) + \mathcal{R}_{img}	70.32	91.30	56.10	94.80	76.45	46.04	47.56	21.30	23.47	25.98
Our T2V-Turbo (MS)	74.47	93.34	58.63	95.80	89.67	45.74	48.47	23.23	25.92	27.51
VCM (VC2)	55.66	63.97	10.81	82.60	79.12	23.06	18.49	25.29	22.31	25.15
VCM (VC2) + \mathcal{R}_{vid}	67.55	87.77	30.38	93.00	86.90	28.81	39.07	25.75	24.65	27.57
VCM (VC2) + \mathcal{R}_{img}	71.70	93.13	46.20	95.00	84.12	37.78	51.34	23.65	24.62	27.75
Our T2V-Turbo (VC2)	74.76	93.96	54.65	95.20	89.90	38.67	55.58	24.42	25.51	28.16

Table 3: Effect of different choices of \mathcal{R}_{vid} . Our T2V-Turbo can always outperform VCM + \mathcal{R}_{img} with either ViCLIP or InternVid2 S2 as \mathcal{R}_{vid} . Table 5 in Appendix E includes further details.

	T2V-Turbo (VC2) $\mathcal{R}_{\text{vid}} = \text{ViCLIP}$	T2V-Turbo (VC2) $\mathcal{R}_{\text{vid}} = \text{InternVid2 S2}$	T2V-Turbo (MS) $\mathcal{R}_{\text{vid}} = \text{ViCLIP}$	T2V-Turbo (MS) $\mathcal{R}_{\text{vid}} = \text{InternVid2 S2}$
Total Score	80.92	81.01	80.62	79.90
Quality Score	82.77	82.57	82.15	82.27
Semantic Score	73.52	74.76	74.47	70.41

4.3 Ablation Studies

We are interested in the effectiveness of each RM, and especially in the impact of the video-text RM \mathcal{R}_{vid} . Therefore, we ablate \mathcal{R}_{img} and \mathcal{R}_{vid} and experiment with different choices of \mathcal{R}_{vid} . In Appendix E, we further experiment with different choices of \mathcal{R}_{img} .

Ablating RMs \mathcal{R}_{img} and \mathcal{R}_{vid} . Recall that the training of our T2V-Turbo incorporate reward feedback from both \mathcal{R}_{img} and \mathcal{R}_{vid} . To demonstrate the effectiveness of each individual RM, we perform ablation study by training VCM (VC2) + \mathcal{R}_{vid} and VCM (VC2) + \mathcal{R}_{img} , which only incorporate feedback from \mathcal{R}_{vid} and \mathcal{R}_{img} , respectively. Again, we evaluate the 4-step generations from different methods on VBench. Results in Table 2 show that incorporating feedback from either \mathcal{R}_{img} or \mathcal{R}_{vid} leads to performance improvement over the baseline VCM. Notably, optimizing \mathcal{R}_{img} alone can already lead to substantial performance gains, while incorporating feedback from \mathcal{R}_{vid} can further improve the Semantic Score on VBench, leading to better text-video alignment. In Appendix H, we qualitatively compare the videos generated by our T2V-Turbo and VCM + \mathcal{R}_{img} , corroborating the effectiveness of our mixture of RMs design.

Effect of different choices of \mathcal{R}_{vid} . We investigate the impact of different choices of \mathcal{R}_{vid} by training T2V-Turbo (VC2) and T2V-Turbo (MS) by setting \mathcal{R}_{vid} as ViCLIP [Wang et al., 2023d] and the second stage model of Intervideo2 (InternVid2 S2). In terms of model architecture, ViCLIP employs the CLIP [Radford et al., 2021] text encoder while InternVid2 S2 leverages the BERT-large [Kenton and Toutanova, 2019] text encoder. Additionally, InternVid2 S2 outperforms ViCLIP in several zero-shot video-text retrieval tasks. As shown in Table 3, T2V-Turbo (VC2) can achieve decent performance on VBench when integrating feedback from either ViCLIP or InternVid2 S2. Conversely, T2V-Turbo (MS) performs better with ViCLIP [Wang et al., 2023d]. Nevertheless, with InternVid2 S2, our T2V-Turbo (MS) still surpasses VCM (MS) + \mathcal{R}_{img} .

5 Related Work

Diffusion-based T2V Models. Many diffusion-based T2V models rely on large-scale image datasets for training [Ho et al., 2022a, Wang et al., 2023c, Chen et al., 2023] or inherit weights from pre-trained text-to-image (T2I) models [Zhang et al., 2023, Blattmann et al., 2023, Khachatryan et al., 2023]. The scale of text-image datasets [Schuhmann et al., 2022] is usually more than ten times the scale of open-sourced video-text datasets [Bain et al., 2021, Wang et al., 2023d] and with higher spatial resolution and diversity [Wang et al., 2023c]. For example, Imagen Video [Ho et al., 2022b] discovers that joint training on a mix of image and video datasets improves the overall visual quality and enables the generation of videos in novel styles. Models trained with WebVid-10M [Bain et al., 2021] like ModelScopeT2V [Wang et al., 2023c] or VideoCrafter [Chen et al., 2023] also treat images as a single-frame video, and use them to improve video qualities. LaVie [Wang et al., 2023b] initialize the training with WebVid-10M and LAION-5B and then continue the training with a curated internal dataset of 23M videos. To overcome the data scarcity of high-quality videos, VideoCrafter2 [Chen et al., 2024] proposes to disentangle motion from appearance at the data level so that it can be trained on high-quality images and low-quality videos. The data limitation of high-quality videos and aligned, accurate video captions has been a longstanding bottleneck of current T2V models. In this paper, we propose to combat this challenge by leveraging reward feedback from a mixture of RMs.

Accelerating inference of Diffusion Models. Various methods have been proposed to accelerate the sampling process of a DM, including advanced numerical ODE solvers [Song et al., 2020b, Lu et al., 2022a,b, Zheng et al., 2022, Dockhorn et al., 2022, Jolicoeur-Martineau et al., 2021] and distillation techniques [Luhman and Luhman, 2021, Salimans and Ho, 2021, Meng et al., 2023, Zheng et al., 2023]. Recently, Consistency Model [Song et al., 2023, Luo et al., 2023a] is proposed to facilitate fast inference by learning a consistency function to map any point at the ODE trajectory to the origin. Li et al. [2024] proposes to augment consistency distillation with an objective to optimize image-text RM to achieve fast and high-quality image generation. Our work extends it for T2V generation, incorporating reward feedback from both an image-text RM and a video-text RM.

Vision-and-language Reward Models. There have been various open-sourced image-text RMs that are trained to mirror human preferences given a text-image pair, including HPS [Wu et al., 2023b,a], ImageReward [Xu et al., 2024a], and PickScore [Kirstain et al., 2024], which are obtained by finetuning a image-text foundation model such as CLIP [Radford et al., 2021] and BLIP [Li et al., 2022], on human preference data. However, to the best of our knowledge, no video-text RMs, e.g., T2VScore [Wu et al., 2024], that mirrors human preference on a text-video pair has been released to the public. In this paper, we choose HPSv2.1 as our image-text RM and directly employ the video foundation models ViCLIP [Wang et al., 2023d] and InterVid S2 [Wang et al., 2024] that are trained for general video-text understanding as our video-text RM. Empirically, we show that incorporating feedback from these RMs can improve the performance of our T2V-Turbo.

Learning from Human/AI Feedback has been proven as an effective way to align the output from a generative model with human preference [Leike et al., 2018, Ziegler et al., 2019, Ouyang et al., 2022, Stiennon et al., 2020, Rafailov et al., 2024, Xu et al., 2024b]. In the field of image generation, various methods have been proposed to align a text-to-image model with human preference, including RL [Sutton and Barto, 2018, Li et al., 2020, 2023a] based methods [Fan et al., 2024, Prabhudesai et al., 2023, Zhang et al., 2024] and backpropagation-based reward finetuning methods [Clark et al., 2023, Xu et al., 2024a, Prabhudesai et al., 2023]. Recently, InstructVideo [Yuan et al., 2023] extends the reward-finetuning methods to optimize a T2V model. However, it still employs an image-text RM to provide reward feedback without considering the transition dynamic of the generated video. In contrast, our work incorporates reward feedback from both an image-text and video-text RM, providing comprehensive feedback to our T2V-Turbo.

6 Conclusion and Limitations

In this paper, we propose T2V-Turbo, achieving both fast and high-quality T2V generation by breaking the quality bottleneck of a VCM. Specifically, we integrate mixed reward feedback into the VCD process of a teacher T2V model. Empirically, we illustrate the applicability of our methods by distilling T2V-Turbo (VC2) and T2V-Turbo (MS) from VideoCrafter2 [Chen et al., 2024] and ModelScopeT2V [Wang et al., 2023c], respectively. Remarkably, the 4-step generations from both our T2V-Turbo outperform SOTA methods on VBench [Huang et al., 2024], even surpassing their

teacher T2V models and proprietary systems including Gen-2 [Esser et al., 2023] and Pika [Pika Labs, 2023]. Our human evaluation further corroborates the results, showing the 4-step generations from our T2V-Turbo are favored by humans over the 50-step DDIM samples from their teacher, which represents over ten-fold inference acceleration with quality improvement.

While our T2V-Turbo marks a critical advancement in efficient T2V synthesis, it is important to recognize certain limitations. Our approach utilizes a mixture of RMs, including a video-text RM \mathcal{R}_{vid} . Due to the lack of an open-sourced video-text RM trained to reflect human preferences on video-text pairs, we instead use video foundation models such as ViCLIP [Wang et al., 2023d] and InternVid S2 [Wang et al., 2024] as our \mathcal{R}_{vid} . Although incorporating feedback from these models has enhanced our T2V-Turbo’s performance, future research should explore the use of a more advanced \mathcal{R}_{vid} for training feedback, which could lead to further performance improvements.

Acknowledgement

The work was funded by an unrestricted gift from Google, and we are grateful for their generous sponsorship. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the sponsors’ official policy, expressed or inferred.

References

- Haoxin Chen, Yong Zhang, Xiaodong Cun, Menghan Xia, Xintao Wang, Chao Weng, and Ying Shan. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. *arXiv preprint arXiv:2401.09047*, 2024.
- Xiang Wang, Shiwei Zhang, Han Zhang, Yu Liu, Yingya Zhang, Changxin Gao, and Nong Sang. Videolcm: Video latent consistency model. *arXiv preprint arXiv:2312.09109*, 2023a.
- 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7346–7356, 2023.
- Pika Labs. Accessed september 25, 2023, 2023. URL <https://www.pika.art/>.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- James Betker, Gabriel Goh, Li Jing, TimBrooks, Jianfeng Wang, Linjie Li, LongOuyang, JuntaoZhuang, JoyceLee, YufeiGuo, WesamManassra, PrafullaDhariwal, CaseyChu, YunxinJiao, and Aditya Ramesh. Improving image generation with better captions. 2023. URL <https://api.semanticscholar.org/CorpusID:264403242>.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. *arXiv preprint arXiv:2403.03206*, 2024.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022.

- Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022.
- Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022a.
- Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. Latent video diffusion models for high-fidelity video generation with arbitrary lengths. *arXiv preprint arXiv:2211.13221*, 2022.
- Yaohui Wang, Xinyuan Chen, Xin Ma, Shangchen Zhou, Ziqi Huang, Yi Wang, Ceyuan Yang, Yanan He, Jiashuo Yu, Peiqing Yang, et al. Lavie: High-quality video generation with cascaded latent diffusion models. *arXiv preprint arXiv:2309.15103*, 2023b.
- David Junhao Zhang, Jay Zhangjie Wu, Jia-Wei Liu, Rui Zhao, Lingmin Ran, Yuchao Gu, Difei Gao, and Mike Zheng Shou. Show-1: Marrying pixel and latent diffusion models for text-to-video generation. *arXiv preprint arXiv:2309.15818*, 2023.
- Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video generation models as world simulators. 2024. URL <https://openai.com/research/video-generation-models-as-world-simulators>.
- Haoxin Chen, Menghan Xia, Yingqing He, Yong Zhang, Xiaodong Cun, Shaoshu Yang, Jinbo Xing, Yaofang Liu, Qifeng Chen, Xintao Wang, et al. Videocrafter1: Open diffusion models for high-quality video generation. *arXiv preprint arXiv:2310.19512*, 2023.
- Jiuniu Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiwei Zhang. Mod-elscope text-to-video technical report. *arXiv preprint arXiv:2308.06571*, 2023c.
- Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *IEEE International Conference on Computer Vision*, 2021.
- Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *International conference on machine learning*, 2023.
- Yang Song and Prafulla Dhariwal. Improved techniques for training consistency models. In *The Twelfth International Conference on Learning Representations*, 2023.
- 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.
- Hangjie Yuan, Shiwei Zhang, Xiang Wang, Yujie Wei, Tao Feng, Yining Pan, Yingya Zhang, Ziwei Liu, Samuel Albanie, and Dong Ni. Instructvideo: Instructing video diffusion models with human feedback. *arXiv preprint arXiv:2312.12490*, 2023.
- Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, SHUM KaShun, and Tong Zhang. Raft: Reward ranked finetuning for generative foundation model alignment. *Transactions on Machine Learning Research*, 2023.
- Kevin Clark, Paul Vicol, Kevin Swersky, and David J Fleet. Directly fine-tuning diffusion models on differentiable rewards. In *The Twelfth International Conference on Learning Representations*, 2023.
- Mihir Prabhudesai, Anirudh Goyal, Deepak Pathak, and Katerina Fragkiadaki. Aligning text-to-image diffusion models with reward backpropagation. *arXiv preprint arXiv:2310.03739*, 2023.
- Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36, 2024a.

- Yaofang Liu, Xiaodong Cun, Xuebo Liu, Xintao Wang, Yong Zhang, Haoxin Chen, Yang Liu, Tiejiong Zeng, Raymond Chan, and Ying Shan. Evalcrafter: Benchmarking and evaluating large video generation models. *arXiv preprint arXiv:2310.11440*, 2023.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2020a.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2020b.
- Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341*, 2023a.
- Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36, 2024.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2021.
- Simian Luo, Yiqin Tan, Suraj Patil, Daniel Gu, Patrick von Platen, Apolinário Passos, Longbo Huang, Jian Li, and Hang Zhao. Lcm-lora: A universal stable-diffusion acceleration module. *arXiv preprint arXiv:2311.05556*, 2023b.
- Yi Wang, Kunchang Li, Xinhao Li, Jiashuo Yu, Yanan He, Guo Chen, Baoqi Pei, Rongkun Zheng, Jilan Xu, Zun Wang, et al. Internvideo2: Scaling video foundation models for multimodal video understanding. *arXiv preprint arXiv:2403.15377*, 2024.
- Yi Wang, Yanan He, Yizhuo Li, Kunchang Li, Jiashuo Yu, Xin Ma, Xinhao Li, Guo Chen, Xinyuan Chen, Yaohui Wang, et al. Internvid: A large-scale video-text dataset for multimodal understanding and generation. *arXiv preprint arXiv:2307.06942*, 2023d.
- Open-Sora. Open-sora: Democratizing efficient video production for all, 2024. URL <https://github.com/hpcaitech/Open-Sora>.
- Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. In *The Eleventh International Conference on Learning Representations*, 2022.
- 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*, pages 8748–8763. PMLR, 2021.
- Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019.
- Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22563–22575, 2023.

- Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15954–15964, 2023.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022.
- Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022b.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022a.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*, 2022b.
- Huangjie Zheng, Pengcheng He, Weizhu Chen, and Mingyuan Zhou. Truncated diffusion probabilistic models and diffusion-based adversarial auto-encoders. *arXiv preprint arXiv:2202.09671*, 2022.
- Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Genie: Higher-order denoising diffusion solvers. *Advances in Neural Information Processing Systems*, 35:30150–30166, 2022.
- Alexia Jolicoeur-Martineau, Ke Li, Rémi Piché-Taillefer, Tal Kachman, and Ioannis Mitliagkas. Gotta go fast when generating data with score-based models. *arXiv preprint arXiv:2105.14080*, 2021.
- Eric Luhman and Troy Luhman. Knowledge distillation in iterative generative models for improved sampling speed. *arXiv preprint arXiv:2101.02388*, 2021.
- Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2021.
- Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and Tim Salimans. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14297–14306, 2023.
- Hongkai Zheng, Weili Nie, Arash Vahdat, Kamyar Aizzadenesheli, and Anima Anandkumar. Fast sampling of diffusion models via operator learning. In *International Conference on Machine Learning*, pages 42390–42402. PMLR, 2023.
- Jiachen Li, Weixi Feng, Wenhui Chen, and William Yang Wang. Reward guided latent consistency distillation. *arXiv preprint arXiv:2403.11027*, 2024.
- Xiaoshi Wu, Keqiang Sun, Feng Zhu, Rui Zhao, and Hongsheng Li. Human preference score: Better aligning text-to-image models with human preference. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2096–2105, 2023b.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022.
- Jay Zhangjie Wu, Guian Fang, Haoning Wu, Xintao Wang, Yixiao Ge, Xiaodong Cun, David Junhao Zhang, Jia-Wei Liu, Yuchao Gu, Rui Zhao, et al. Towards a better metric for text-to-video generation. *arXiv preprint arXiv:2401.07781*, 2024.
- Jan Leike, David Krueger, Tom Everitt, Miljan Martic, Vishal Maini, and Shane Legg. Scalable agent alignment via reward modeling: a research direction. *arXiv preprint arXiv:1811.07871*, 2018.

- Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*, 2019.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744, 2022.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021, 2020.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36, 2024.
- Wenda Xu, Jiachen Li, William Yang Wang, and Lei Li. Bpo: Supercharging online preference learning by adhering to the proximity of behavior LLM. In *The 2024 Conference on Empirical Methods in Natural Language Processing*, 2024b.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Jiachen Li, Quan Vuong, Shuang Liu, Minghua Liu, Kamil Ciosek, Henrik Christensen, and Hao Su. Multi-task batch reinforcement learning with metric learning. *Advances in neural information processing systems*, 33:6197–6210, 2020.
- Jiachen Li, Edwin Zhang, Ming Yin, Qinxun Bai, Yu-Xiang Wang, and William Yang Wang. Offline reinforcement learning with closed-form policy improvement operators. In *International Conference on Machine Learning*, pages 20485–20528. PMLR, 2023a.
- Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. Reinforcement learning for fine-tuning text-to-image diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Yinan Zhang, Eric Tzeng, Yilun Du, and Dmitry Kislyuk. Large-scale reinforcement learning for diffusion models. *arXiv preprint arXiv:2401.12244*, 2024.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021.
- Zhen Li, Zuo-Liang Zhu, Ling-Hao Han, Qibin Hou, Chun-Le Guo, and Ming-Ming Cheng. Amt: All-pairs multi-field transforms for efficient frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9801–9810, 2023b.
- Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 402–419. Springer, 2020.
- Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5148–5157, 2021.
- Jialian Wu, Jianfeng Wang, Zhengyuan Yang, Zhe Gan, Zicheng Liu, Junsong Yuan, and Lijuan Wang. Grit: A generative region-to-text transformer for object understanding. *arXiv preprint arXiv:2212.00280*, 2022.
- Kunchang Li, Yali Wang, Yizhuo Li, Yi Wang, Yinan He, Limin Wang, and Yu Qiao. Unmasked teacher: Towards training-efficient video foundation models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19948–19960, 2023c.

Kaiyi Huang, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation. *Advances in Neural Information Processing Systems*, 36:78723–78747, 2023a.

Xinyu Huang, Youcai Zhang, Jinyu Ma, Weiwei Tian, Rui Feng, Yuejie Zhang, Yaqian Li, Yandong Guo, and Lei Zhang. Tag2text: Guiding vision-language model via image tagging. In *The Twelfth International Conference on Learning Representations*, 2023b.

Appendix

A Experiment and Hyperparameter (HP) Details

When performing qualitative comparisons between different methods, we ensure to use the same random seed for head-to-head video comparisons.

As mentioned in Sec. 4, we train T2V-Turbo (VC2) and T2V-Turbo (MS) by distilling from the teacher diffusion-based T2V models VideoCrafter2 [Chen et al., 2024] and the less advanced ModelScopeT2V [Wang et al., 2023c], respectively. Specifically, VideoCrafter2 supports video FPS as an input and generates videos at the resolution of 512x320. For simplicity, we always set FPS to 16 when distilling our T2V-Turbo (VC2) from VideoCrafter2. On the other hand, ModelScopeT2V always generates video at 8 FPS at a resolution of 256x256.

We conduct our training with the WebVid10M [Bain et al., 2021] datasets. Note that both teacher T2V models are also trained with WebVid10M. We train our models for 10K gradient steps with 6 - 8 NVIDIA A100 GPUs without gradient accumulation and set the batch size to 1 for each GPU device. That is, we load 1 video with 16 frames. At each training iteration, we always sample 16 frames from the input video. We employ HPSv2.1 [Wu et al., 2023a] as our image-text RM \mathcal{R}_{img} . When distilling from VideoCrafter2, we utilize the 2nd Stage model of InternVideo2 [Wang et al., 2024] as our video-text RM \mathcal{R}_{vid} . When distilling from ModelScopeT2V, we set \mathcal{R}_{vid} to be ViCLIP [Wang et al., 2023d]. To optimize J_{img} (8), we randomly sample 6 frames from the video by setting $M = 6$. For the hyperparameters, we set learning rate $1e - 5$ and guidance scale range $[\omega_{\min}, \omega_{\max}] = [5, 15]$. We use DDIM [Song et al., 2020b] as our ODE solver Ψ and set the skipping step $k = 20$. For T2V-Turbo (VC2), we set $\beta_{\text{img}} = 1$ and $\beta_{\text{vid}} = 2$. For T2V-Turbo (MS), we set $\beta_{\text{img}} = 2$ and $\beta_{\text{vid}} = 3$.

As mentioned in Sec. 2, we employ the DDIM [Song et al., 2020b] ODE solver Ψ_{DDIM} by following the practice of Luo et al. [2023a]. Its formula from t_{n+k} to t_n is given below

$$\Psi_{\text{DDIM}}(z_{t_{n+k}}, t_{n+k}, t_n, \mathbf{c}) = \underbrace{\frac{\alpha_{t_n}}{\alpha_{t_{n+k}}} z_{t_{n+k}} - \beta_{t_n} \left(\frac{\beta_{t_{n+k}} \cdot \alpha_{t_n}}{\alpha_{t_{n+k}} \cdot \beta_{t_n}} - 1 \right) \hat{\epsilon}_{\psi}(z_{t_{n+k}}, \mathbf{c}, t_{n+k})}_{\text{DDIM Estimated } z_{t_n}} - z_{t_{n+k}} \quad (11)$$

where $\hat{\epsilon}_{\psi}$ denotes the noise prediction model from the teacher T2V model. We refer interested readers to the original LCM paper [Luo et al., 2023a] for further details.

B Pseudocodes for Training our T2V-Turbo

We include the pseudocodes for training our T2V-Turbo in Algorithm 1. We use the red color to highlight the difference from the standard (latent) consistency distillation [Luo et al., 2023a, Song et al., 2023].

Algorithm 1 T2V-Turbo Training Pipeline

Require: text-video dataset \mathcal{D} , initial model parameter θ , learning rate η , ODE solver Ψ , distance metric d , EMA rate μ , noise schedule $\alpha(t), \beta(t)$, guidance scale $[\omega_{\min}, \omega_{\max}]$, skipping interval k , VAE encoder \mathcal{E} , **decoder \mathcal{D} , image-text RM \mathcal{R}_{img} , video-text RM \mathcal{R}_{vid} , reward scale β_{img} and β_{vid} .**
Encoding training data into latent space: $\mathcal{D}_z = \{(\mathbf{z}, \mathbf{c}) \mid \mathbf{z} = E(\mathbf{x}), (\mathbf{x}, \mathbf{c}) \in \mathcal{D}\}$
 $\theta^- \leftarrow \theta$
repeat
 Sample $(\mathbf{z}, \mathbf{c}) \sim \mathcal{D}_z, n \sim \mathcal{U}[1, N - k]$ and $\omega \sim [\omega_{\min}, \omega_{\max}]$
 Sample $\mathbf{z}_{t_{n+k}} \sim \mathcal{N}(\alpha(t_{n+k}) \mathbf{z}; \sigma^2(t_{n+k}) \mathbf{I})$
 $\hat{\mathbf{z}}_{t_n}^{\Psi, \omega} \leftarrow \mathbf{z}_{t_{n+k}} + (1 + \omega) \Psi(\mathbf{z}_{t_{n+k}}, t_{n+k}, t_n, \mathbf{c}) - \omega \Psi(\mathbf{z}_{t_{n+k}}, t_{n+k}, t_n, \emptyset)$
 $\hat{\mathbf{x}}_0 = \mathcal{D}(\mathbf{f}_\theta(\mathbf{z}_{t_{n+k}}, \omega, \mathbf{c}, t_{n+k}))$
 $J_{\text{img}}(\theta) = \mathbb{E}_{\hat{\mathbf{x}}_0, \mathbf{c}} \left[\sum_{m=1}^M \mathcal{R}_{\text{img}}(\hat{\mathbf{x}}_0^m, \mathbf{c}) \right]$
 $J_{\text{vid}}(\theta) = \mathbb{E}_{\hat{\mathbf{x}}_0, \mathbf{c}} [\mathcal{R}_{\text{vid}}(\hat{\mathbf{x}}_0, \mathbf{c})]$
 $L_{\text{CD}} = d(\mathbf{f}_\theta(\mathbf{z}_{t_{n+k}}, \omega, \mathbf{c}, t_{n+k}), \mathbf{f}_{\theta^-}(\hat{\mathbf{z}}_{t_n}^{\Psi, \omega}, \omega, \mathbf{c}, t_n))$
 $\mathcal{L}(\theta, \theta^-; \Psi) \leftarrow L_{\text{CD}} - \beta_{\text{img}} J_{\text{img}}(\theta) - \beta_{\text{vid}} J_{\text{vid}}(\theta)$
 $\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}(\theta, \theta^-)$
 $\theta^- \leftarrow \text{stop_grad}(\mu \theta^- + (1 - \mu) \theta)$
until convergence

C Further Details about VBench

We provide a brief introduction of the metrics included in VBench [Huang et al., 2024] followed by introducing the derivation rules for the **Quality Score**, **Semantic Score** and **Total Score**. We refer interested readers to read the VBench paper for further details.

The following metrics are used to construct the **Quality Score**.

- **Subject Consistency** (Subject Consist.) is calculated by the DINO [Caron et al., 2021] feature similarity across video frames.
- **Background Consistency** (BG Consist.) is calculated by CLIP [Radford et al., 2021] feature similarity across video frames.
- **Temporal Flickering** (Temporal Flicker.) is computed by the mean absolute difference across video frames.
- **Motion Smoothness** (Motion Smooth.) is evaluated by motion priors in the video frame interpolation model [Li et al., 2023b].
- **Aesthetic Quality** is calculated by mean of aesthetic scores evaluated by the LAION aesthetic predictor [Schuhmann et al., 2022].
- **Dynamic Degree** is calculated using RAFT [Teed and Deng, 2020].
- **Image Quality** is evaluated by the MUSIQ [Ke et al., 2021] image quality predictor.

Quality Score is calculated as the weighted sum of the normalized scores of each metric mentioned above. The weight for all metrics is 1, except for **Dynamic Degree**, which has a weight of 0.5.

The following metrics are used to construct the **Semantic Score**.

- **Object Class** is calculated by detecting the success rate of generating the object specified by the user using GRiT [Wu et al., 2022].
- **Multiple Object** is calculated by detecting the success rate of generating all objects specified in the prompt using GRiT [Wu et al., 2022].
- **Human Action** is evaluated by the UMT model [Li et al., 2023c].
- **Color** is calculated by comparing the color caption generated by GRiT [Wu et al., 2022] against the expected color.
- **Spatial Relationship** (Spatial Relation.) is calculated by a rule-based method similar to [Huang et al., 2023a].
- **Scene** is calculated by comparing the video captions generated by Tag2Text [Huang et al., 2023b] against the scene descriptions in the prompt.
- **Appearance Style** (Appear Style.) is calculated by using ViCLIP [Wang et al., 2023d] to compare the video feature and the style description in the user prompt.
- **Temporal Style** is calculated based on the similarity between the video feature and the style description feature provided by ViCLIP [Wang et al., 2023d].
- **Overall Consistency** (Overall Consist.) is calculated based on the similarity between the video feature and the entire text prompt feature provided by ViCLIP [Wang et al., 2023d]. ViCLIP [Wang et al., 2023d]

Semantic Score is simply calculated as the mean of the normalized scores of each metric mentioned above. And the **Total Score** is the weighted sum of **Quality Score** and **Semantic Score**, which is given by

$$\text{Total Score} = \frac{4 \cdot \text{Quality Score} + \text{Total Score}}{5} \quad (12)$$

Table 4: **Automatic Evaluation on VBench** [Huang et al., 2024]. We compare our T2V-Turbo (VC2) and T2V-Turbo (MS) with baseline methods across the 16 VBench dimensions. A higher score indicates better performance for a particular dimension. We bold the best results for each dimension and underline the second-best result. **Quality Score** is calculated with the 7 dimensions from the top table. **Semantic Score** is calculated with the 9 dimensions from the bottom table. **Total Score** a weighted sum of **Quality Score** and **Semantic Score**. Both our T2V-Turbo (VC2) and T2V-Turbo (MS) **surpass all baseline methods with 4 inference steps** in terms of Total Score, including the proprietary systems Gen-2 and Pika.

Models	Total Score	Quality Score	Subject Consist.	BG Consist.	Temporal Flicker.	Motion Smooth.	Aesthetic Quality	Dynamic Degree	Image Quality
CogVideo	67.01	72.06	92.19	96.20	97.64	96.47	38.18	42.22	41.03
VideoCrafter0.9	73.02	74.91	86.24	92.88	97.60	91.79	44.41	89.72	57.22
ModelScopeT2V	75.75	78.05	89.87	95.29	98.28	95.79	52.06	<u>66.39</u>	58.57
Open-Sora	75.91	78.82	92.09	97.39	98.41	95.61	57.76	48.61	61.51
LaVie	77.08	78.78	91.41	97.47	98.30	96.38	54.94	49.72	61.90
LaVie-Interpolation	77.12	79.07	92.00	97.33	98.78	97.82	54.00	46.11	59.78
Show-1	78.93	80.42	95.53	98.02	99.12	98.24	57.35	44.44	58.66
VideoCrafter1	79.72	81.59	95.10	98.04	98.93	95.67	62.67	55.00	65.46
Pika	80.40	82.68	96.76	98.95	99.77	<u>99.51</u>	<u>63.15</u>	37.22	62.33
VideoCrafter2	80.44	82.20	<u>96.85</u>	98.22	98.41	97.73	<u>63.13</u>	42.50	67.22
Gen-2	80.58	82.47	97.61	97.61	<u>99.56</u>	99.58	66.96	18.89	67.42
VCM (MS)	75.84	78.80	93.06	97.30	98.51	98.00	48.99	46.11	61.98
Our T2V-Turbo (MS)	<u>80.62</u>	82.15	94.82	<u>98.71</u>	97.99	95.64	60.04	<u>66.39</u>	<u>68.09</u>
VCM (VC2)	73.97	78.54	94.02	96.05	99.06	98.84	54.56	42.50	52.72
Our T2V-Turbo (VC2)	81.01	<u>82.57</u>	96.28	97.02	97.48	97.34	63.04	49.17	72.49

Models	Semantic Score	Object Class	Multiple Objects	Human Action	Color	Spatial Relation.	Scene	Appear. Style	Temporal Style	Overall Consist.
CogVideo	46.83	73.40	18.11	78.20	79.57	18.24	28.24	22.01	7.80	7.70
VideoCrafter0.9	65.46	87.34	25.93	93.00	78.84	36.74	43.36	21.57	25.42	25.21
ModelScopeT2V	66.54	82.25	38.98	92.40	81.72	33.68	39.26	23.39	25.37	25.67
Open-Sora	64.28	74.98	33.64	85.00	78.15	43.95	37.33	21.58	25.46	26.18
LaVie	70.31	91.82	33.32	96.80	86.39	34.09	52.69	23.56	<u>25.93</u>	26.41
LaVie-Interpolation	69.31	90.68	30.93	<u>95.80</u>	85.69	30.06	52.62	23.53	26.01	26.51
Show-1	72.98	93.07	45.47	95.60	86.35	53.50	47.03	23.06	25.28	27.46
VideoCrafter1	72.22	78.18	45.66	91.60	93.32	58.86	43.75	24.41	25.54	26.76
Pika	71.26	87.45	46.69	88.00	85.31	<u>65.65</u>	44.80	21.89	24.44	25.47
VideoCrafter2	73.42	92.55	40.66	95.00	<u>92.92</u>	<u>35.86</u>	<u>55.29</u>	<u>25.13</u>	25.84	28.23
Gen-2	73.03	90.92	<u>55.47</u>	89.20	89.49	66.91	48.91	19.34	24.12	26.17
VCM (MS)	63.98	83.18	24.85	87.20	85.72	31.57	42.44	23.20	23.30	24.18
Our T2V-Turbo (MS)	<u>74.47</u>	<u>93.34</u>	58.63	95.80	89.67	45.74	48.47	23.23	25.92	27.51
VCM (VC2)	55.66	63.97	10.81	82.60	79.12	23.06	18.49	25.29	22.31	25.15
Our T2V-Turbo (VC2)	74.76	93.96	54.65	95.20	89.90	38.67	55.58	24.42	25.51	<u>28.16</u>

D Human Evaluation Details

Figure 5 shows the user interface displayed to the labelers when conducting our human evaluations. Each method generate videos of 16 frames using the 700 prompts from EvalCrafter [Liu et al., 2023]. For our T2V-Turbo (VC2), we collect its 4-step and 8-step generations and compare them with the 50-step DDIM samples from its teacher VideoCrafter2. For our T2V-Turbo (MS), we collect its 2-step and 4-step generations and compare them with the 50-step DDIM samples from its teacher ModelScopeT2V. We also compare the 4-step generations between our T2V-Turbo their baseline VCM, demonstrating the significant quality improvement of our methods.

As mentioned in Sec. 4.2, we hire labelers from Amazon Mechanical Turk platform and form the video comparison task as many batches of HITs. Specifically, we choose labelers from English-speaking countries, including AU, CA, NZ, GB, and the US. Each task needs around 30 seconds to complete, and we pay each submitted HIT with 0.2 US dollars. Therefore, the hourly payment is about 24 US dollars.

We note that the data annotation part of our project is classified as exempt by Human Subject Committee via IRB protocols.

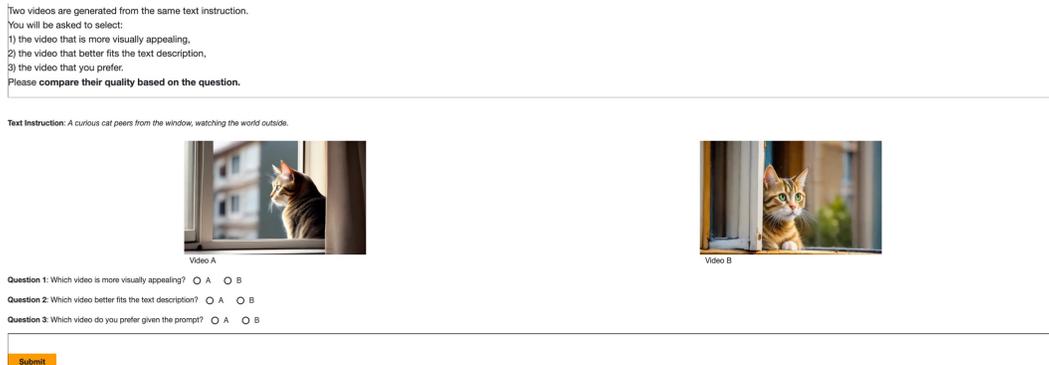


Figure 5: User interface of our human evaluation experiments.

E Additional Ablation Studies

In this section, we provide the full ablation results performed in Table 3, which can be found in Table 5. We further examine the effect of different choices of \mathcal{R}_{img} . In the initial stage of our project, we train VCM (VC2) + \mathcal{R}_{img} with several different image-text RMs, including HPSv2.1 [Wu et al., 2023a], PickScore [Kirstain et al., 2024], and ImageReward [Xu et al., 2024a]. We collect the 4-step generations from each method and qualitatively compare them with the 4-step generation from the baseline VCM. As shown in Figure 6, incorporating reward feedback from any of these \mathcal{R}_{img} leads to quality improvement over the baseline VCM (VC2). It is worth noting that HPSv2.1 and PickScore are fine-tuned from CLIP with human preference data. Therefore, learning from CLIP might also lead to better performance than the baseline VCM.

Table 5: Effect of different choices of \mathcal{R}_{vid} . T2V-Turbo (VC2) can achieve decent performance on VBench when integrating feedback from either ViCLIP or InternVid2 S2. On the other hand, T2V-Turbo (MS) achieves a better result with ViCLIP [Wang et al., 2023d].

Models	Total Score	Quality Score	Subject Consist.	BG Consist.	Temporal Flicker.	Motion Smooth.	Aesthetic Quality	Dynamic Degree	Image Quality
T2V-Turbo (MS), $\mathcal{R}_{\text{vid}} = \text{ViCLIP}$	80.62	82.15	94.82	98.71	97.99	95.64	60.04	66.39	68.09
T2V-Turbo (MS), $\mathcal{R}_{\text{vid}} = \text{InternVid2 S2}$	79.90	82.27	96.68	99.36	97.74	95.66	65.30	52.22	68.23
T2V-Turbo (VC2), $\mathcal{R}_{\text{vid}} = \text{ViCLIP}$	80.92	82.77	96.93	97.47	98.03	97.48	63.38	43.61	72.94
T2V-Turbo (VC2), $\mathcal{R}_{\text{vid}} = \text{InternVid2 S2}$	81.01	82.57	96.28	97.02	97.48	97.34	63.04	49.17	72.49

Models	Semantic Score	Object Class	Multiple Objects	Human Action	Color	Spatial Relation.	Scene	Appear. Style	Temporal Style	Overall Consist.
T2V-Turbo (MS), $\mathcal{R}_{\text{vid}} = \text{ViCLIP}$	74.47	93.34	58.63	95.80	89.67	45.74	48.47	23.23	25.92	27.51
T2V-Turbo (MS), $\mathcal{R}_{\text{vid}} = \text{InternVid2 S2}$	70.41	94.05	48.73	92.60	81.69	45.41	48.15	21.45	23.84	26.24
T2V-Turbo (VC2), $\mathcal{R}_{\text{vid}} = \text{ViCLIP}$	73.52	94.05	50.52	94.40	89.85	36.77	54.17	23.81	25.34	28.11
T2V-Turbo (VC2), $\mathcal{R}_{\text{vid}} = \text{InternVid2 S2}$	74.76	93.96	54.65	95.20	89.90	38.67	55.58	24.42	25.51	28.16

F Qualitative Results

We provide additional qualitative comparisons between our T2V-Turbo, the baseline VCM, and their teacher T2V models in Figures 7, 8, 9, and 10.

The prompts for the top two and bottom two rows in Figure 1 are given below:

- With the style of low-poly game art, A majestic, white horse gallops gracefully across a moonlit beach.
- Kung Fu Panda posing in cyberpunk, neonpunk style.

G Broader Impact

The ability to create highly realistic synthetic videos raises concerns about misinformation and deepfakes, which can be used to manipulate public opinion, defame individuals, or perpetrate fraud. Addressing these concerns requires robust regulatory frameworks and ethical guidelines to ensure the technology is used responsibly and for the benefit of society. **Therefore, we are committed to installing safeguard when releasing our models.** Specifically, we will require users to adhere to usage guidelines.

Despite the challenges, the impact of our T2V-Turbo is profound, offering a scalable solution that significantly enhances the accessibility and practicality of generating high-quality videos at a remarkable speed. This innovation not only broadens the potential applications in fields ranging from digital art to visual content creation but also sets a new benchmark for future research in T2V synthesis, emphasizing the importance of human-centric design in the development of generative AI technologies.

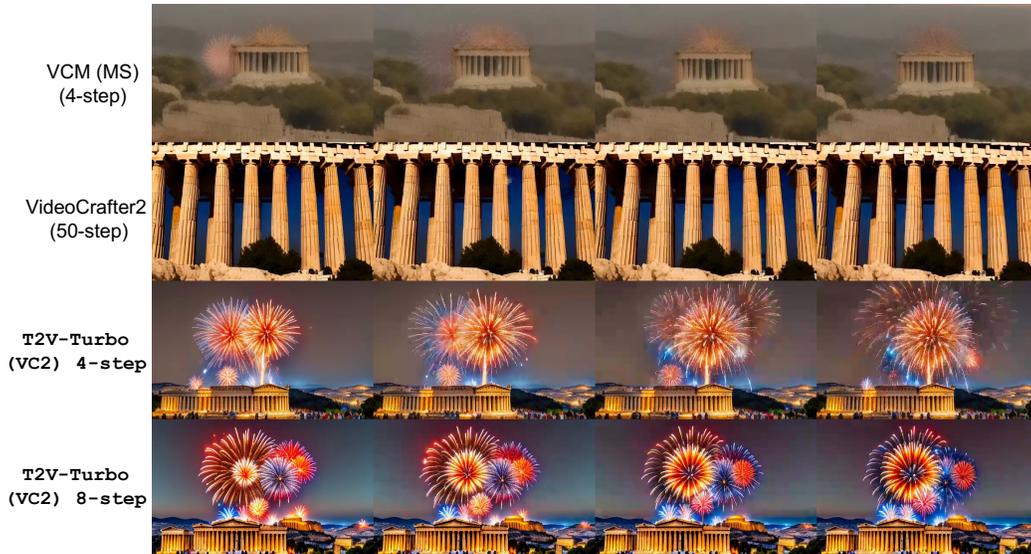


Figure 6: Ablation study on the choice of the \mathcal{R}_{img} . We compare the 4-step generations from each methods. The three \mathcal{R}_{img} we tested can all improve the video generation quality compare to the baseline VCM (VC2).

A wise tortoise in a tweed hat and spectacles reads a newspaper, Howard Hodgkin style



drove viewpoint, fireworks above the Parthenon



Iron Man is walking towards the camera in the rain at night, with a lot of fog behind him. Science fiction movie, close-up

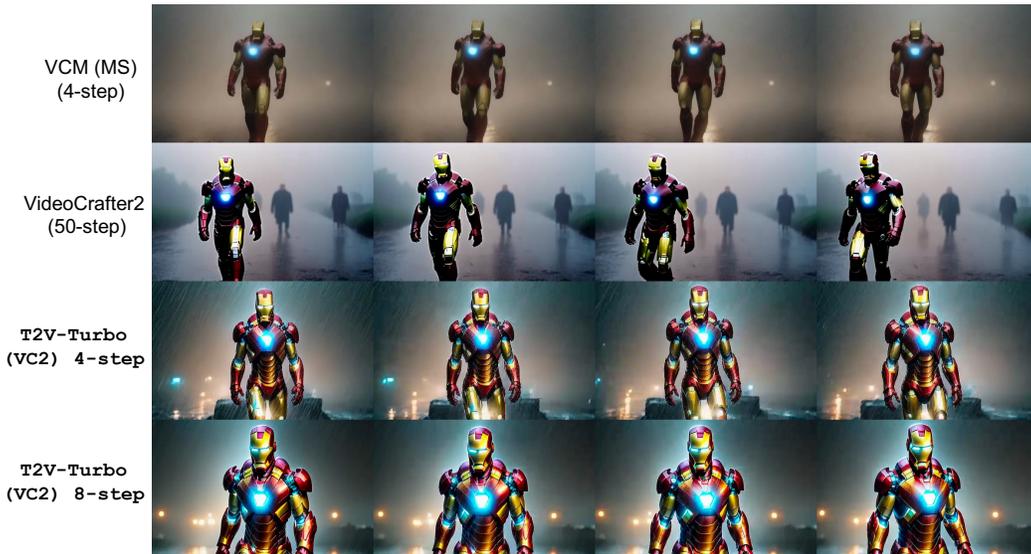
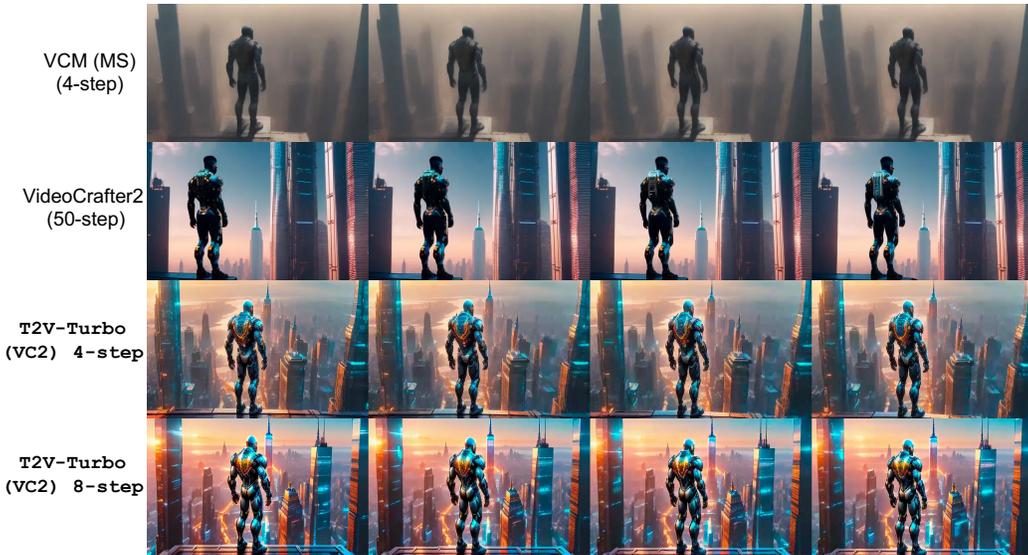


Figure 7: Additional qualitative comparison results for our T2V-Turbo (VC2).

With the style of sketch, A sophisticated monkey in a beret and striped shirt paints in a French artist's studio.



a cyborg standing on top of a skyscraper, overseeing the city, back view, cyberpunk vibe, 2077, NYC, intricate details, 4K



A Egyptian tomp hieroglyphics painting of A regal lion, decked out in a jeweled crown, surveys his kingdom.

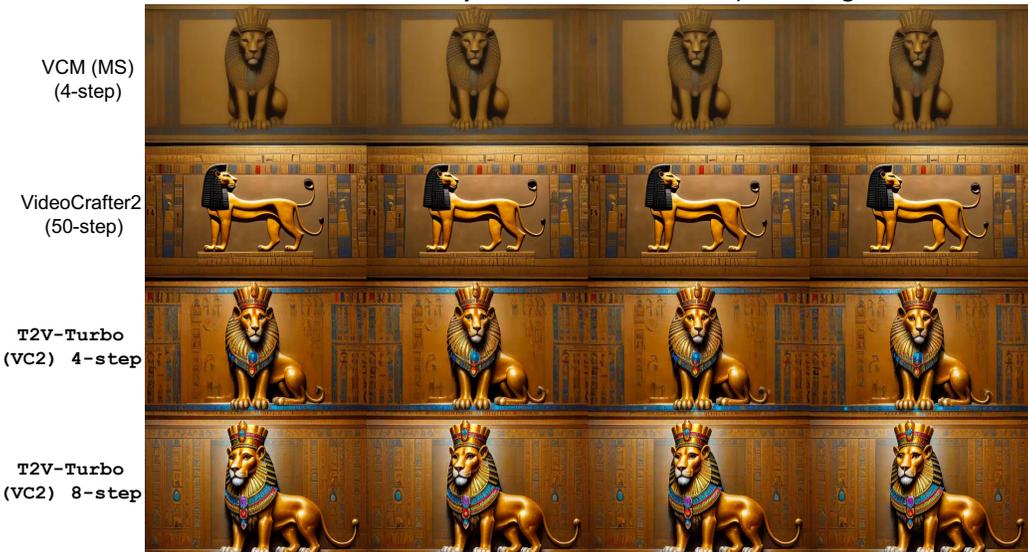
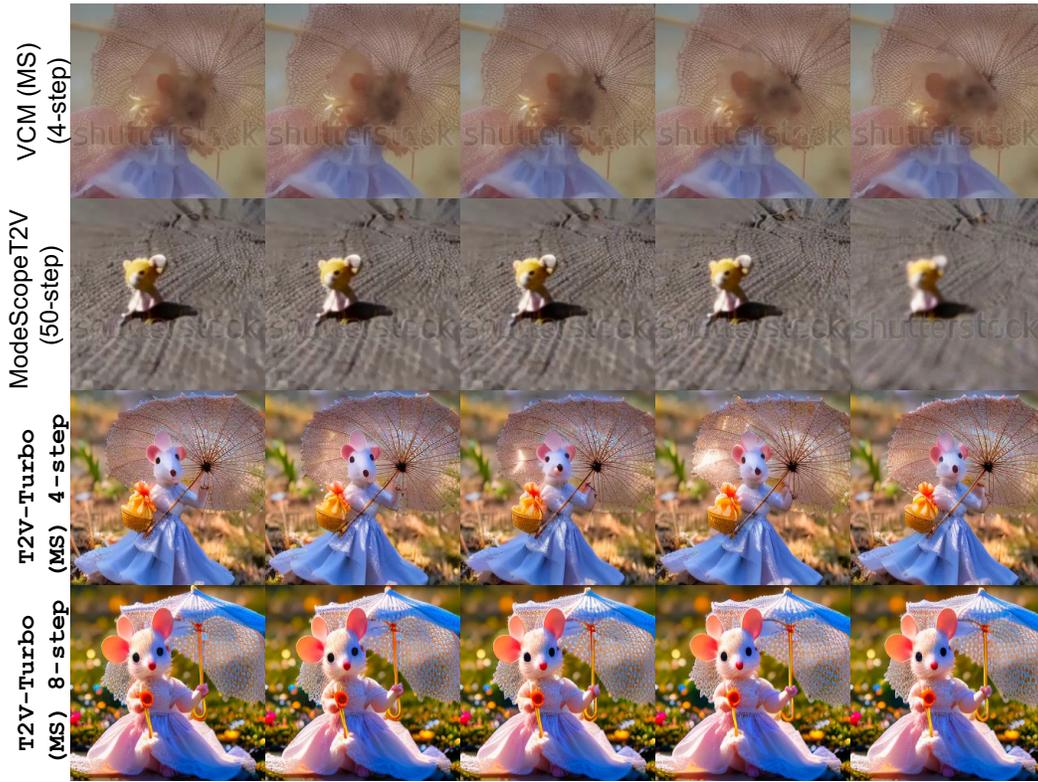


Figure 8: Additional qualitative comparison results for our T2V-Turbo (VC2).

Macro len style, A tiny mouse in a dainty dress holds a parasol to shield from the sun.



pop art style, A photo of an astronaut riding a horse in the forest. There is a river in front of them with water lilies.



Figure 9: Additional qualitative comparison results for our T2V-Turbo (MS).

Mickey Mouse is dancing on white background



a man looking at a distant mountain in Sci-fi style

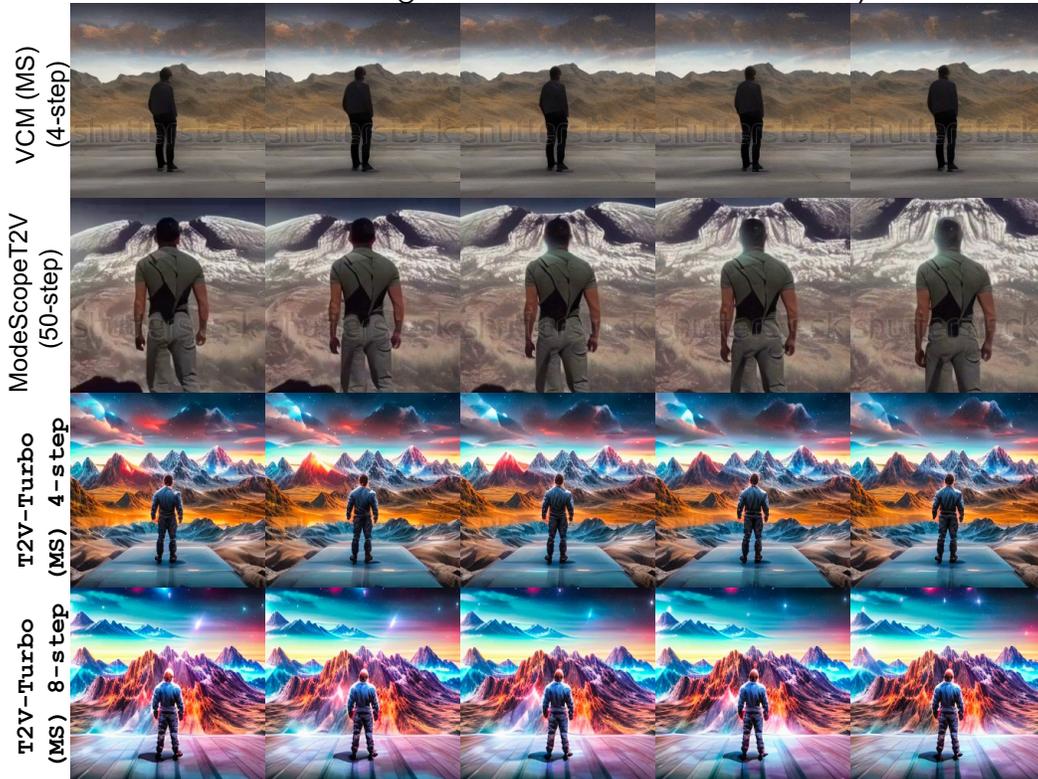


Figure 10: Additional qualitative comparison results for our T2V-Turbo (MS).

H Comparing videos generated by T2V-Turbo and VCM + \mathcal{R}_{img}

Please click to play videos in **Adobe Acrobat**.

Prompt: *A panda standing on a surfboard in the ocean in sunset.*

T2V-Turbo (VC2)

VCM (VC2) + \mathcal{R}_{img}

Analysis The panda on the **right** is instead **sitting** on the surfboard.

Prompt: *A raccoon is playing the electronic guitar.*

T2V-Turbo (VC2)

VCM (VC2) + \mathcal{R}_{img}

Analysis The **right** video fails on **playing the electronic guitar**.

Prompt: *A motorcycle accelerating to gain speed.*

T2V-Turbo (VC2)

VCM (VC2) + \mathcal{R}_{img}

Analysis The motorcycle on the **right** is actually moving backward.

Prompt: *A squirrel eating a burger.*

T2V-Turbo (VC2)

VCM (VC2) + \mathcal{R}_{img}

Analysis The squirrel on the **right** looks more like it is **holding** a burger.

Prompt: *A Mars rover moving on Mars.*

T2V-Turbo (VC2)

VCM (VC2) + \mathcal{R}_{img}

Analysis The hills on the **right** in the background also move.

Prompt: *A horse galloping across an open field.*

T2V-Turbo (VC2)

VCM (VC2) + \mathcal{R}_{img}

Analysis Another horse suddenly runs into the scene of the **right** video.

Prompt: *A black vase.*

T2V-Turbo (VC2)

VCM (VC2) + \mathcal{R}_{img}

Analysis The **right** video shows **two vases** instead of one.

Prompt: *Happy dog wearing a yellow turtleneck, studio, portrait, dark background.*

T2V-Turbo (VC2)

VCM (VC2) + \mathcal{R}_{img}

Analysis The dog on the **right** doesn't look happy.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We have made sure to accurately illustrate our main claims in the abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We address the limitation of our work in Sec. 6

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when the image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in the appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We have described our experimental settings in detail in Sec. 4 and Appendix A. We will also release the codes and models.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We include our codes in the supplementary materials. We have released our models and codes in <https://github.com/Ji4chenLi/t2v-turbo>.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide experimental details in both Sec. 4 and Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in the appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We only have the computational resources to run the training for one time. And thus do not include error bar.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We include the information in Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We conform, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss the broader impact of our work in Sec. G.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [\[Yes\]](#)

Justification: We have discussed this in Appendix G.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [\[Yes\]](#)

Justification: We have make sure that we have followed all the rules.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We have provided detailed usage guidance.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [Yes]

Justification: We performance human evaluation on our methods in Sec. 4.2. We include further details including full text of instructions given to participants and screenshots in Appendix D.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [Yes]

Justification: The data annotation part of the project is classified as exempt by Human Subject Committee via IRB protocols.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.