

VideoCrafter2: Overcoming Data Limitations for High-Quality Video Diffusion Models

Haoxin Chen Yong Zhang* Xiaodong Cun Menghan Xia
 Xintao Wang Chao Weng Ying Shan

Tencent AI Lab

In cyberpunk, neonpunk style, Kung Fu Panda, jump and kick. *Cinematic photo melting pistachio ice cream dripping down the cone. 35mm photograph, film, bokeh.* *Large motion, surrounded by butterflies, a girl walks through a lush garden.*

Figure 1. Give a text prompt, our method can generate a video with high visual quality and accurate text-video alignment. Note that it is trained with only low-quality videos and high-quality images. No high-quality videos are required. *Best viewed with Acrobat Reader. Click the images to play the video clips.*

Abstract

Text-to-video generation aims to produce a video based on a given prompt. Recently, several commercial video models have been able to generate plausible videos with minimal noise, excellent details, and high aesthetic scores. However, these models rely on large-scale, well-filtered, high-quality videos that are not accessible to the community. Many existing research works, which train models using the low-quality WebVid-10M dataset, struggle to generate high-quality videos because the models are optimized to fit WebVid-10M. In this work, we explore the training scheme of video models extended from Stable Diffusion and investigate the feasibility of leveraging low-quality videos and synthesized high-quality images to obtain a high-quality video model. We first analyze the connection between the spatial

and temporal modules of video models and the distribution shift to low-quality videos. We observe that full training of all modules results in a stronger coupling between spatial and temporal modules than only training temporal modules. Based on this stronger coupling, we shift the distribution to higher quality without motion degradation by finetuning spatial modules with high-quality images, resulting in a generic high-quality video model. Evaluations are conducted to demonstrate the superiority of the proposed method, particularly in picture quality, motion, and concept composition.

1. Introduction

Benefiting from the development of diffusion models [25, 43], video generation has achieved breakthroughs, particularly in basic text-to-video (T2V) generation models. Most existing methods [14, 23, 26, 30, 48, 63] follow a logic to

* Corresponding author.

obtain video models, *i.e.*, extending a text-to-image (T2I) backbone to a video model by adding temporal modules and then training it with videos. Several models train video models from scratch, while most start from a pre-trained T2I model, typically Stable Diffusion (SD) [38]. Models can also be categorized into two groups based on the space modeled by diffusion models, *i.e.*, pixel-space models [26, 30, 60] and latent-space models [14, 23, 48, 63]. The latter is the dominant approach. Picture quality, motion consistency, and concept composition are essential dimensions for evaluating a video model.

Recently, a few commercial startups have released their T2V models [5, 6, 8, 9] that can produce plausible videos with minimal noise, excellent details, and high aesthetic scores. However, they are trained on a large-scale and well-filtered high-quality video dataset, which is not accessible to the community and academia. Collecting millions of high-quality videos is challenging due to copyright restrictions and post-filtering processing. Though there are a few open-source video datasets collected from the Internet for video understanding, such as HowTo100M [33], HD-VILA-100M [56], and InterVid [51], there exist many issues for video generation, *e.g.*, poor picture quality and caption, multiple clips in one video, and static frames or slides. WebVid-10M [12] is the most widely used dataset to train video generation models in academia. The clips are well-segmented, and the diversity is good. However, the picture quality is unsatisfactory, and most videos have a resolution of about 320p. The lack of high-quality datasets poses a significant obstacle to training high-quality video models in academia.

In this work, we target a quite challenging problem, *i.e.*, training high-quality video models without using high-quality videos. We dive into the training process of SD-based video models to analyze the connection between spatial and temporal modules under different training strategies and investigate the distribution shift to low-quality videos. We make a meaningful observation that the full training of all modules results in a stronger coupling between appearance and motion than just training temporal modules. The full training can achieve more natural motion and tolerate more subsequent modification of spatial modules, which is the key to improving the quality of generated videos. Based on the observation of the connection, we propose a method to overcome the data limitation by disentangling motion from appearance at the data level. Specifically, instead of high-quality videos, we exploit low-quality videos to guarantee motion consistency and use high-quality images to ensure picture quality and concept composition ability. Benefiting from the successful T2I models such as SDXL and Midjourney, it is convenient to obtain a large set of images with high-resolution and complex concept composition. As shown in Fig. 2, we design a pipeline to train a video diffusion model following the guidelines from the analysis. Fig. 1

shows visual examples generated by our method.

Our main contributions are summarized as follows:

- We propose a method to overcome the data for training high-quality video models by disentangling motion from appearance at the data level.
- We investigate the connection between spatial and temporal modules, and the distribution shift. We identify the keys to obtain a high-quality video model.
- We design an effective pipeline based on the observations, *i.e.*, obtaining a fully trained video model first and tuning the spatial modules with synthesized high-quality images.

2. Related Work

The evolution of video generation techniques goes along with the development of generative models. Generative adversarial networks [17] and variational auto-encoders [18] are the commonly used backbones in early research of video generation, *e.g.*, VGAN [47], TGAN [40], MoCoGAN [44], GODIA [52], StyleGAN-V [41], and MCVD [46]. Then, since transformers have been successfully applied in various fields, they are also introduced for video synthesis, *e.g.*, CogVideo [28], VideoGPT [57], NUVA-infinity [53], TATS [19], MAGVIT [58], Phenaki [45].

Recently, diffusion models (DMs) [25, 42, 43] have been a famous star in generative models, especially in text-to-image (T2I) generation [13, 21, 24, 34, 36–39, 62]. For video generation, Video Diffusion Models (VDMs) are proposed to model the distribution of videos. VDM [27] is the first to utilize a space-time factorized U-Net to model videos in pixel space for unconditional video generation. It uses an image-video joint training strategy to avoid concept forgetting. Imagen Video [26] and Make-a-Video [30] are two cascade models that target text-to-video generation in pixel space. Show-1 [60] is another cascade model that uses IF [1] as the base model and LDM extended video model for super-resolution. Subsequently, LVDM [15, 23] and MagicVideo [63] propose to extend LDM [38] to model videos in the latent space of an auto-encoder. Many other methods use the same paradigm, including ModelScope [48], Align Your Latent [14], Hotshot-XL [7], LAVIE [50], PY-OCO [20], VideoFactory [49], VPDM [59], VIDM [32], and Latent-Shift [11]. Besides text-to-video generation, a few methods, such as [16, 55, 61], attempt to generate videos from a given image and a prompt as condition.

Several startups release their text-to-video generation services, *e.g.*, Gen-2 [5], Pika Labs [9], Moonvalley [8], and Genmo [6]. Their models can generate plausible results with minimal noise, excellent details, and high aesthetic scores. However, those methods are trained with a large-scale well-filtered high-quality video dataset that is not accessible to researchers. The video models are also not available, leading to the slow development of downstream tasks to a certain extent. The most widely used video dataset is WebVid-10M,

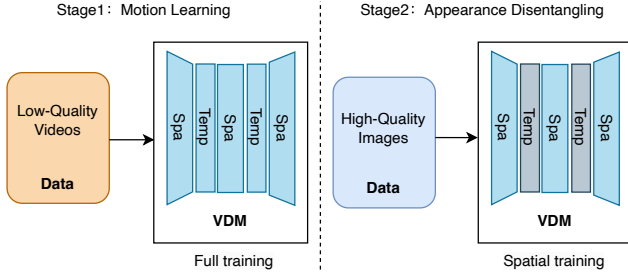


Figure 2. The efficient training pipeline. In the first stage, we use a large amount of low-quality videos to fully train the VDM to learn motion. In the second stage, we only train the spatial parameters with the high-quality images generated by the T2I model to improve the concept composition and picture quality.

a large-scale dataset of short videos with textual descriptions sourced from stock footage sites. The videos are diverse and rich in their content, and each video is well-segmented, however, the picture quality is unsatisfactory and most videos are 320p. Training a high-quality video model under the data limitation is quite challenging.

AnimateDiff [22] finds that combining temporal modules from a video model trained on WebVid-10M and a LORA SD model can improve the picture quality of the generated videos. However, this is not a generic model and does not always work. There are a few severe issues. First, the temporal modules can only be combined with a few selected LORA models, which makes it not a generic model. Second, the motion quality degenerates when the modules do not match the LORA model well. Unlike AnimateDiff, we analyze the connection between spatial and temporal modules instead of direct combination, and design a pipeline to train a generic high-quality video model without high-quality video by disentangling appearance and motion at the data level.

3. Method

We propose an effective method to overcome the data limitation for training high-quality video diffusion models. We first analyze the connection between spatial and temporal modules of SD-based video models under different training strategies. Based on the observations, we then develop a pipeline to train high-quality video models with just low-quality videos and high-quality images, *i.e.*, disentangling appearance from motion at the data level.

3.1. Spatial-temporal Connection Analyses

Base T2V model. To leverage the prior in SD trained on a large-scale image dataset, most text-to-video diffusion models inflate the SD model to a video model by adding temporal modules, including Align Your Latent [14], AnimateDiff [22], LVDM [23], Magic Video [63], ModelScopeT2V [48], and LAVIE [50]. They follow VDM [27]

to use a particular type of 3D-UNet that is factorized over space and time.

These models can be categorized into two groups according to their training strategies. One is to use videos to train temporal modules with spatial ones fixed, called *partial training*. The other is to learn both the spatial and temporal modules with the SD weights as initialization, called *full training*.

Though these SD-based T2V models have similar architectures, they are trained under different training settings. We use one typical model to investigate the connection between spatial and temporal modules under the two training strategies. We follow the architecture of the open-source VideoCrafter1 [15] with FPS (frames per second) condition. We also incorporate the temporal convolution in ModelScopeT2V [48] to improve temporal consistency.

Parameter Perturbation for Full and Partial Training.

We apply the two training strategies to the same architecture using the same data. The model is initialized from pretrained SD weights. WebVid-10M [12] is exploited as the training data. To avoid concept forgetting, LAION-COCO 600M [3] is also used for video and image joint training. The resolution is 512×320 . For simplicity, the fully trained video model is denoted as $M_F(\theta_T, \theta_S)$, while the partially trained one is denoted as $M_P(\theta_T, \theta_S^0)$, where θ_T and θ_S are the learned parameters of the temporal and spatial modules, respectively. θ_S^0 are the original spatial parameters of SD.

To evaluate the connection strength between spatial and temporal modules, we perturb the parameters of the specified modules by using another high-quality image dataset \mathcal{D}_I for finetuning. The image data is JDB [35] that consists of synthesized images from Midjourney [4]. As the JDB has 4 million images, we use LORA [29] for finetuning.

Spatial Perturbation. We first perturb the spatial parameters of the two video models using the image dataset. The temporal parameters are frozen. The perturbation process of the fully trained base model M_F can be denoted as:

$$M'_F(\theta_T, \theta_S + \Delta_{\theta_S}) \leftarrow \text{PERTB}_{\theta_S}^{\text{LORA}}(M_F(\theta_T, \theta_S), \mathcal{D}_I),$$

where $\text{PERTB}_{\theta_S}^{\text{LORA}}$ denote finetuning M_F with respect to θ_S on the image dataset \mathcal{D}_I using LORA. Δ_{θ_S} represents the parameters of the LORA branch. Similarly, we can obtain the perturbed model of the partially trained video model:

$$M'_P(\theta_T, \theta_S^0 + \Delta_{\theta_S}) \leftarrow \text{PERTB}_{\theta_S}^{\text{LORA}}(M_P(\theta_T, \theta_S^0), \mathcal{D}_I).$$

For easy understanding, we also use the name ‘F-Spa-LORA’ to denote model M'_F and ‘P-Spa-LORA’ for M'_P . ‘F’ denotes the fully trained base model while ‘P’ stands for the partially trained model. ‘Spa’ and ‘Temp’ mean finetuning spatial and temporal modules, respectively. ‘LORA’ represents using LORA for finetuning, while ‘DIR’ means direct

iter = 1K

iter = 10K

P-Temp-LORA

F-Temp-LORA

P-Spa-LORA

F-Spa-LORA

A professional dancer gracefully performs ballet on stage.

Figure 3. Perturbing spatial modules using LORA. Best viewed with Acrobat Reader. Click the images to play the video clips.

finetuning without LORA. For example, ‘F-Spatial-LORA’ represents perturbing spatial modules of the fully trained T2V model using LORA.

Comparing the synthesized videos of the two resulting models, we have the following observations. First, the motion quality of F-Spa-LORA is more stable than that of P-Spa-LORA (see user study in Table 4). The motion of P-Spa-LORA becomes worse quickly during the finetuning process. The more finetuning steps, the video tends to be more still with local flicker (see Fig. 3). While the motion of F-Spa-LORA slightly degenerates compared to the fully trained base model. Second, P-Spa-LORA achieves much better visual quality than F-Spa-LORA (see Fig. 3). The picture quality and aesthetic score of F-Spa-LORA are greatly improved compared to the partially trained base model (see Table 3). Surprisingly, the watermark is also removed. While F-Spa-LORA obtains a slight improvement in picture quality and aesthetic score, the generated videos are still noisy.

From the two observations, we can conclude that the coupling strength between spatial and temporal modules of the fully trained model is stronger than that of the partially trained model. Because the spatial-temporal coupling of the partially trained model can be easily broken, leading to quick motion degeneration and picture quality shift. A stronger connection can tolerate parameter perturbation more than a weak one. Our observation can be used to explain the quality improvement and motion degeneration of AnimateDiff. AnimateDiff is not a generic model and only works for selected personalized SD models. The reason is that its motion modules are obtained with the partially training strategy, and

Robot dancing in times square.

Figure 4. Perturbing temporal modules using LORA. Best viewed with Acrobat Reader. Click the images to play the video clips.

they cannot tolerate large parameter perturbations. When the personalized model does not match the temporal modules, both picture and motion quality will degenerate.

Temporal Perturbation. The partially trained model has only the temporal modules updated, but the picture quality is shifted to the quality of WebVid-10M. Hence, the temporal modules take responsibility for not only the motion but also the picture quality. We perturb the temporal modules while fixing the spatial modules with the image dataset. The perturbation processes can be denoted as:

$$M_F''(\theta_T + \Delta\theta_T, \theta_S) \leftarrow \text{PERTB}_{\theta_T}^{\text{LORA}}(M_F(\theta_T, \theta_S), \mathcal{D}_I),$$
$$M_P''(\theta_T + \Delta\theta_T, \theta_S^0) \leftarrow \text{PERTB}_{\theta_T}^{\text{LORA}}(M_P(\theta_T, \theta_S^0), \mathcal{D}_I).$$

We observe that the picture quality of P-Temp-LORA (M_P'') is better than F-Temp-LORA (M_F''). However, the foreground and background of the videos are more shaky, *i.e.*, the temporal consistency becomes worse (see Fig. 4). The picture of F-Temp-LORA is improved, but the watermark is still there. Its motion is close to the base model and much better than P-Temp-LORA (see Table 4). Those observations also support the conclusion obtained from spatial perturbation.

3.2. Data-level Disentanglement of Appearance and Motion

Since obtaining a large-scale, high-quality video dataset with high diversity is challenging due to copyright issues, we explore the possibility of training a high-quality video model without using high-quality videos. We have access to low-quality videos such as WebVid-10M and high-quality images such as JDB. We propose to disentangle motion from appearance at the data level, *i.e.*, learning motion from low-quality videos while learning picture quality and aesthetics from high-quality images. We can first train a video model with videos and then fine-tune the video model with images. *The keys lie in how to train a video model and how to fine-tune it with images.*

According to the study of the connection between spatial and temporal modules, a fully trained model is more suitable

An astronaut is waving his hands on the moon.

Figure 5. Module selection based on the fully trained T2V model. *Best viewed with Acrobat Reader. Click the images to play the video clips.*

Anime illustration of a blue pig, riding a scooter near a lake, with the sun in the sky

F-Spa-DIR-LAION

F-Spa-DIR

Figure 6. Influence of image data on concept composition. ‘F-Spa-DIR-LAION’ uses the LAION aesthetics V2 as the image data while ‘F-Spa-DIR’ uses JDB. *Best viewed with Acrobat Reader. Click the images to play the video clips.*

for the subsequent finetuning with high-quality images. This is because the strong spatial-temporal coupling can tolerate the parameter perturbation for both spatial and temporal modules without obvious motion degeneration.

Next, we need to investigate how to fine-tune the base model with images. In both spatial and temporal perturbation (Sec. 3.1), the picture quality can be improved but not very significantly. To obtain a greater quality improvement, we evaluate two strategies. One is to involve more parameters, *i.e.*, finetuning both spatial and temporal modules with images. The other is to change the finetuning method, *i.e.*, using direct finetuning without LORA. We can evaluate the following four cases:

$$\begin{aligned} M_F^A(\theta_T + \Delta\theta_T, \theta_S + \Delta\theta_S) &\leftarrow \text{PERTB}_{\theta_T, \theta_S}^{\text{LORA}}(M_F(\theta_T, \theta_S), \mathcal{D}_I), \\ M_F^B(\theta_T, \theta_S + \Delta\theta_S) &\leftarrow \text{PERTB}_{\theta_S}(M_F(\theta_T, \theta_S), \mathcal{D}_I), \\ M_F^C(\theta_T + \Delta\theta_T, \theta_S) &\leftarrow \text{PERTB}_{\theta_T}(M_F(\theta_T, \theta_S), \mathcal{D}_I), \\ M_F^D(\theta_T + \Delta\theta_T, \theta_S + \Delta\theta_S) &\leftarrow \text{PERTB}_{\theta_T, \theta_S}(M_F(\theta_T, \theta_S), \mathcal{D}_I), \end{aligned}$$

Where M_F^A (F-Spa&Temp-LORA) is obtained by following the first strategy, M_F^B , M_F^C , and M_F^D are obtained using

the second strategy. M_F^B (F-Spa-DIR) and M_F^C (F-Temp-DIR) represent directly finetuning the spatial and temporal modules, respectively. M_F^D (F-Spa&Temp-DIR) represents directly finetuning all modules.

Comparing the generated videos of the four models, we have the following observations. First, F-Spa&Temp-LORA can further improve the picture quality of F-Spa-LORA, but the quality is still unsatisfying. The watermark exists in most generated videos, and the noise is obvious. Second, F-Temp-DIR achieves better picture quality than F-Temp-LORA. It is also better than F-Spa&Temp-LORA. The watermark is removed or lightened in half of the videos. Third, F-Spa-DIR and F-Spa&Temp-DIR achieve the best picture quality among the fine-tuned models. However, the motion of F-Spa-DIR is better (see Fig. 5 and Table 4). The foreground and background of F-Spa&Temp-DIR are flashing in videos generated by M_F^D , especially local textures.

By exploring the finetuning strategies and different modules, we identify that directly finetuning spatial modules with high-quality images is the best way to improve the picture quality without marginal loss of motion quality. As shown in Fig. 2, our data-level disentanglement pipeline can be summarized as follows: fully training a video model with low-quality videos first and then directly finetuning the spatial modules only with high-quality images.

3.3. Promotion of Concept Composition

To improve the concept composition ability of video models, we propose to use synthesized images with complex concepts instead of using real images at the partial finetuning stage. The success of T2I models such as SDXL and Midjourney is built upon large-scale high-quality images. They have the ability to composite concepts that do not appear in the real world. Rather than using their training images, we propose transferring their concept composition ability to video models by synthesizing a set of images with complex concepts. In this way, we can alleviate the burden of capturing both concept and motion well at the same time.

To validate the effectiveness of synthesized images, we use JDB and LAION-aesthetics V2 [2] as image data for the second finetuning stage. LAION-aesthetics V2 consists of web-collected images while JDB contains images synthesized by Midjourney. We observe that the model trained with JDB has much better concept composition ability (see Fig. 6 and Table 3). More results are in the supplementary material.

4. Experiments

4.1. Settings

Data. To overcome data limitations, we utilized WebVid-10M [12] as the source of low-quality video data and JDB [35] for high-quality image data. WebVid-10M is a large-scale, diverse video dataset comprising approximately 10 million text-video pairs. The resolution of most videos is 336×596 , and each video consists of a single shot. During training, we sample from videos at varying frame rates. JDB is a large-scale image dataset featuring around 4 million high-resolution images from Midjourney, each annotated with a corresponding text prompt. To prevent concept forgetting during the training of the base T2V model, we also employ LAION-COCO[3], a dataset comprising 600 million generated high-quality captions for publicly available web images, for both image and video training.

Metrics. We exploit EvalCrafter [31] for quantitative evaluation. EvalCrafter is a benchmark to evaluate text-to-video generation models, which contains around 18 objective metrics for visual quality, content quality, motion quality, and text-caption alignment. It provides about 512 prompts. The objective metrics are aligned to user opinions from five subjective studies, i.e., motion quality, text-video alignment, temporal consistency, visual quality, and user favor. The motion quality considers three metrics: action recognition, average flow, amplitude classification score, while temporal consistency considers warping error, semantic consistency, face consistency. The technical and aesthetic scores in EvalCrafter are adapted from DOVER [54]. Besides, we conduct user studies of human preference since there still lacks a comprehensive objective metric to measure motion quality.

Training Details. In Sec 3.1, the two based models share the same architecture, adapted from the open-source VideoCrafter1 [15], and incorporate temporal convolution from ModelScopeT2V [48]. The spatial modules are initialized with weights from SD 2.1, and the outputs of the temporal modules are initialized to zeros. The training resolution is set at 512×320 . For joint image and video training, we utilize the low-quality WebVid-10M and LAION-COCO datasets. The models are trained on 32 NVIDIA A100 GPUs for 270K iterations with a batch size of 128. The learning rate is set at 5×10^{-5} for all training tasks. When employing

	Visual Quality	Text-Video Alignment	Motion Quality	Temporal Consistency
PikaLab*	63.52	54.11	57.74	69.35
Gen2*	67.35	52.30	62.53	69.71
VideoCrafter1	61.64	66.76	56.06	60.36
Show-1	52.19	62.07	53.74	60.83
AnimateDiff	58.89	74.79	51.38	56.61
Ours	63.28	64.67	53.95	62.02

Table 1. Comparison on the EvalCrafter benchmark. Higher score indicates better performance. * commercial models.

LORA for the perturbation of temporal or spatial modules, we exclusively use JDB for tuning. The finetuning is conducted on 8 A100 GPUs for 30K iterations with a batch size of 256. Given that the images from JDB have a square resolution, we adjust the finetuning resolution to 512×512 .

4.2. Comparison with State-of-the-Art T2V Models

We compare our approach with several state-of-the-art T2V models, including popular commercial models such as Gen-2 [5] and Pika Labs [9], as well as open-source models like Show-1 [60], VideoCrafter1 [15], and AnimateDiff [22]. Gen-2, Pika Labs, and VideoCrafter1 all utilize high-quality videos for training their T2V models. It is noteworthy that AnimateDiff and our models use only the videos from WebVid-10M. Show-1 employs additional high-quality videos for finetuning to eliminate the watermark in WebVid-10M. AnimateDiff is not a generic T2V model; it works only when the LORA SD model is compatible with its temporal modules. For our comparison, we use its temporal modules (second version) based on SD v1.5 and employ Realistic Vision V2.0 [10] as its corresponding LORA model.

Quantitative Evaluation. The quantitative results obtained using EvalCrafter are presented in Table 1. Our method achieves visual quality comparable to that of VideoCrafter1 and Pika Labs, which use high-quality videos for training. This underscores the effectiveness of employing high-quality images to enhance picture quality and aesthetic scores. Furthermore, our text-video alignment performance is ranked second. In terms of motion quality, our performance surpasses that of Show-1 but falls short of models that utilize a larger volume of videos to learn motion. This indicates that our method can enhance visual quality without significant motion degradation.

Qualitative Evaluation. The visual comparison is depicted in Fig. 7. Additional results are provided in the supplementary material. The visual quality of our results is on par with that of commercial models such as Gen-2 and Pika Labs. Since we employ JDB as the image dataset, the picture quality of our synthesized videos shifts from WebVid-10M to JDB. Regarding motion, our motion quality is superior to

A bear rummages through a dumpster, searching for food scraps *A Time 1980 painting of a boy going to school* *A man cruises through the city on a motorcycle, feeling the adrenaline rush* *A monkey eating a pizza in central park, GoPro film style*

Gen-2
Pika Labs
VideoCrafter1
Show-1
AnimateDiff
Ours

Figure 7. Comparison of different text-to-video generation models. *Best viewed with Acrobat Reader. Click the images to play the video clips.*

that of AnimateDiff and comparable to Show-1. Although the integration of temporal modules with a LORA SD model can enhance visual quality, AnimateDiff experiences motion

degradation in generic scenes.

Methods	Text-Video Alignment	Motion Quality	Visual Quality
Ours vs Gen2	56%	46%	34%
Ours vs AnimateDiff	55%	64%	69%
Ours vs Show-1	59%	59%	82%
Ours vs VideoCrafter1	61%	63%	61%

Table 2. Human preference. The numbers represent the probability of users choosing our method.

User Study. For further evaluation, we conduct a user study to compare our method with other video models. We select 50 prompts from EvalCrafter, covering diverse scenes, styles, and objects. When comparing a model pair, three video production experts are asked to select their preferred video from three options: method 1, method 2, and comparable results, according to the given subject, *i.e.*, visual quality, motion quality, and text-video alignment. The results are shown in Table 4. Our method has better visual quality than AnimateDiff and Show-1 and is comparable to VideoCrafter1. Our method is more preferred than Show-1 and AnimateDiff in motion quality.

4.3. Strategy Evaluation

Spatial-temporal Connection. In Sec. 3.1, we show a visual comparison of perturbing the spatial and temporal parameters of the fully and partially trained models in Fig. 3 and Fig. 4. Here we provide the quantitative comparisons about the visual quality in Table 3, including aesthetic and technical scores from DOVER [54]. We observe that finetuning the partially trained model can always achieve better visual quality than the fully trained model. It means that the distribution of the partially trained model can be shifted more easily. Besides, we conduct a user study asking participants to choose a favorable model that has better performance in motion, in terms of foreground/background flash and motion flicker. The results are shown in Table 4. It can be observed that the motion quality of perturbed fully trained models is better. The fully trained model can tolerate larger parameter perturbations than the partially trained model. These observations show that the fully trained model has stronger spatial-temporal coupling.

Module Selection. After selecting the fully trained model as the base, we use two strategies to identify the most effective module to fine-tune, resulting in four models in Sec. 3.2. The visual quality evaluation of these models is shown in the bottom part of Table 3. The visual quality of F-Spa-DIR and F-Spa&Temp-DIR is much better than the other two models. It reveals that directly finetuning spatial modules is the key to improving picture quality.

Since F-Spa-DIR and F-Spa&Temp-DIR achieve close visual quality, we conduct a user study on motion quality to

Method	Aesthetic Score (\uparrow)	Technical Score (\uparrow)
P-base	34.32	42.69
F-base	46.55	51.76
P-Spa-LORA	78.25	72.74
F-Spa-LORA	77.97	59.60
P-Temp-LORA	77.40	54.85
F-Temp-LORA	66.26	50.32
F-Spa-DIR	82.57	70.35
F-Temp-DIR	82.77	65.34
F-Spa&Temp-DIR	83.59	67.75
F-Spa&Temp-LORA	80.44	63.61
F-Spa-DIR-LAION	67.83	54.26

Table 3. Visual quality evaluation of the perturbed T2V models.

Methods	Motion Quality
F-Spa-LORA vs P-Spa-LORA	87%
F-Temp-LORA vs P-Temp-LORA	73%
F-Spa-DIR vs F-Spa&Temp-DIR	67%

Table 4. User study on the motion of the perturbed T2V models.

determine the final model. The results are shown in the last row of Table 4. Directly finetuning the spatial modules only performs better in motion. As shown in Fig. 5, F-Spa-DIR is more stable and has better temporal consistency than F-Spa&Temp-DIR. The latter has obvious flashes in both the foreground and background.

Influence of Image Data. To verify the effectiveness of synthesized images, we use the LAION Aesthetics V2 dataset and JDB to directly fine-tune the spatial modules in the second stage, respectively. The visual examples are shown in Fig. 6. It shows that the model trained with JDB composite concepts better than the model trained with LAION Aesthetics V2. The quantitative evaluation of visual quality is shown in Table 3. F-Spa-DIR is much better than F-Spa-DIR-LAION in both aesthetic and technical scores.

5. Conclusion

To overcome data limitations, we propose a method for training high-quality video diffusion models without using high-quality videos. We delve into the training schemes of SD-based video models and investigate the coupling strength between spatial and temporal dimensions. We observe that fully trained T2V models exhibit stronger spatial-temporal coupling than partially trained models. Based on this observation, we propose disentangling appearance from motion at the data level, *i.e.*, by exploiting low-quality videos for motion learning and high-quality images for appearance learning. Additionally, we suggest using synthetic images with complex concepts for finetuning, rather than real images. Quantitative and qualitative evaluations are conducted to demonstrate the effectiveness of the proposed method.

References

- [1] If. Accessed October 22, 2023 [Online] <https://github.com/deep-floyd/IF>. 2
- [2] Laion-aesthetics. Accessed October 22, 2023 [Online] <https://laion.ai/blog/laion-aesthetics/>, . 6
- [3] Laion-coco. Accessed October 22, 2023 [Online] <https://laion.ai/blog/laion-coco/>, . 3, 6
- [4] Midjourney. Accessed October 22, 2023 [Online] <https://www.midjourney.com/home>. 3
- [5] Gen-2. Accessed October 22, 2023 [Online] <https://research.runwayml.com/gen2>, . 2, 6
- [6] Genmo. Accessed October 22, 2023 [Online] <https://www.genmo.ai/>, . 2
- [7] Hotshot-xl. Accessed October 22, 2023 [Online] <https://github.com/hotshotco/Hotshot-XL>. 2
- [8] Moonvalley. Accessed October 22, 2023 [Online] <https://moonvalley.ai/>. 2
- [9] Pika labs. Accessed October 22, 2023 [Online] <https://www.pika.art/>. 2, 6
- [10] Realistic vision v2.0. Accessed October 22, 2023 [Online] <https://huggingface.co/ckpt/realistic-vision-v2.0>. 6
- [11] Jie An, Songyang Zhang, Harry Yang, Sonal Gupta, Jia-Bin Huang, Jiebo Luo, and Xi Yin. Latent-shift: Latent diffusion with temporal shift for efficient text-to-video generation. *arXiv preprint arXiv:2304.08477*, 2023. 2
- [12] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021. 2, 3, 6
- [13] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022. 2
- [14] 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 *CVPR*, 2023. 1, 2, 3
- [15] 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. 2, 3, 6
- [16] Xinyuan Chen, Yaohui Wang, Lingjun Zhang, Shaobin Zhuang, Xin Ma, Jiashuo Yu, Yali Wang, Dahua Lin, Yu Qiao, and Ziwei Liu. Seine: Short-to-long video diffusion model for generative transition and prediction. *arXiv preprint arXiv:2310.20700*, 2023. 2
- [17] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35(1):53–65, 2018. 2
- [18] Carl Doersch. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*, 2016. 2
- [19] Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and Devi Parikh. Long video generation with time-agnostic vqgan and time-sensitive transformer. In *ECCV*, pages 102–118. Springer, 2022. 2
- [20] Songwei Ge, Seungjun Nah, Guilin Liu, Tyler Poon, Andrew Tao, Bryan Catanzaro, David Jacobs, Jia-Bin Huang, Ming-Yu Liu, and Yogesh Balaji. Preserve your own correlation: A noise prior for video diffusion models. In *ICCV*, 2023. 2
- [21] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *CVPR*, 2022. 2
- [22] Yuwei Guo, Ceyuan Yang, Anyi Rao, Yaohui Wang, Yu Qiao, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*, 2023. 3, 6
- [23] 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. 1, 2, 3
- [24] Yingqing He, Shaoshu Yang, Haoxin Chen, Xiaodong Cun, Menghan Xia, Yong Zhang, Xintao Wang, Ran He, Qifeng Chen, and Ying Shan. Scalecrafter: Tuning-free higher-resolution visual generation with diffusion models. *arXiv preprint arXiv:2310.07702*, 2023. 2
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 1, 2
- [26] 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*, 2022. 1, 2
- [27] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. In *NeurIPS*, 2022. 2, 3
- [28] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*, 2022. 2
- [29] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 3
- [30] Yaosi Hu, Chong Luo, and Zhenzhong Chen. Make it move: controllable image-to-video generation with text descriptions. In *CVPR*, 2022. 1, 2
- [31] Yaofang Liu, Xiaodong Cun, Xuebo Liu, Xintao Wang, Yong Zhang, Haoxin Chen, Yang Liu, Tiejong Zeng, Raymond Chan, and Ying Shan. Evalcrafter: Benchmarking and evaluating large video generation models, 2023. 6
- [32] Kangfu Mei and Vishal Patel. Vidm: Video implicit diffusion models. In *AAAI*, pages 9117–9125, 2023. 2
- [33] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *ICCV*, 2019. 2

- [34] Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *ICML*, 2022. 2
- [35] Junting Pan, Keqiang Sun, Yuying Ge, Hao Li, Haodong Duan, Xiaoshi Wu, Renrui Zhang, Aojun Zhou, Zipeng Qin, Yi Wang, et al. Journeydb: A benchmark for generative image understanding. *arXiv preprint arXiv:2307.00716*, 2023. 3, 6
- [36] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 2
- [37] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2
- [39] 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. *NeurIPS*, 2022. 2
- [40] Masaki Saito, Eiichi Matsumoto, and Shunta Saito. Temporal generative adversarial nets with singular value clipping. In *ICCV*, pages 2830–2839, 2017. 2
- [41] Ivan Skorokhodov, Sergey Tulyakov, and Mohamed Elhoseiny. Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. In *CVPR*, pages 3626–3636, 2022. 2
- [42] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. 2
- [43] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021. 1, 2
- [44] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *CVPR*, pages 1526–1535, 2018. 2
- [45] Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang, Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. Phenaki: Variable length video generation from open domain textual description. *arXiv preprint arXiv:2210.02399*, 2022. 2
- [46] Vikram Voleti, Alexia Jolicoeur-Martineau, and Chris Pal. Mcvd-masked conditional video diffusion for prediction, generation, and interpolation. In *NeurIPS*, 2022. 2
- [47] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. *NeurIPS*, 29, 2016. 2
- [48] Jiuniu Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiwei Zhang. Modelscope text-to-video technical report. *arXiv preprint arXiv:2308.06571*, 2023. 1, 2, 3, 6
- [49] Wenjing Wang, Huan Yang, Zixi Tuo, Huiguo He, Junchen Zhu, Jianlong Fu, and Jiaying Liu. Videofactory: Swap attention in spatiotemporal diffusions for text-to-video generation. *arXiv preprint arXiv:2305.10874*, 2023. 2
- [50] Yaohui Wang, Xinyuan Chen, Xin Ma, Shangchen Zhou, Ziqi Huang, Yi Wang, Ceyuan Yang, Yinan He, Jiashuo Yu, Peiqing Yang, et al. Lavie: High-quality video generation with cascaded latent diffusion models. *arXiv preprint arXiv:2309.15103*, 2023. 2, 3
- [51] Yi Wang, Yinan He, Yizhuo Li, Kunchang Li, Jiashuo Yu, Xin Ma, Xinyuan Chen, Yaohui Wang, Ping Luo, Ziwei Liu, et al. Internvid: A large-scale video-text dataset for multimodal understanding and generation. *arXiv preprint arXiv:2307.06942*, 2023. 2
- [52] Chenfei Wu, Lun Huang, Qianxi Zhang, Binyang Li, Lei Ji, Fan Yang, Guillermo Sapiro, and Nan Duan. Godiva: Generating open-domain videos from natural descriptions. *arXiv preprint arXiv:2104.14806*, 2021. 2
- [53] Chenfei Wu, Jian Liang, Xiaowei Hu, Zhe Gan, Jianfeng Wang, Lijuan Wang, Zicheng Liu, Yuejian Fang, and Nan Duan. Nuwa-infinity: Autoregressive over autoregressive generation for infinite visual synthesis. *arXiv preprint arXiv:2207.09814*, 2022. 2
- [54] Haoning Wu, Erli Zhang, Liang Liao, Chaofeng Chen, Jingwen Hou, Annan Wang, Wenxiu Sun, Qiong Yan, and Weisi Lin. Exploring video quality assessment on user generated contents from aesthetic and technical perspectives. In *CVPR*, pages 20144–20154, 2023. 6, 8
- [55] Jinbo Xing, Menghan Xia, Yong Zhang, Haoxin Chen, Xintao Wang, Tien-Tsin Wong, and Ying Shan. Dynamicrafter: Animating open-domain images with video diffusion priors. *arXiv preprint arXiv:2310.12190*, 2023. 2
- [56] Hongwei Xue, Tiankai Hang, Yanhong Zeng, Yuchong Sun, Bei Liu, Huan Yang, Jianlong Fu, and Baining Guo. Advancing high-resolution video-language representation with large-scale video transcriptions. In *CVPR*, 2022. 2
- [57] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021. 2
- [58] Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, et al. Magvit: Masked generative video transformer. In *CVPR*, 2023. 2
- [59] Sihyun Yu, Kihyuk Sohn, Subin Kim, and Jinwoo Shin. Video probabilistic diffusion models in projected latent space. In *CVPR*, pages 18456–18466, 2023. 2
- [60] 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, 2023. 2, 6
- [61] Shiwei Zhang, Jiayu Wang, Yingya Zhang, Kang Zhao, Hangjie Yuan, Zhiwu Qin, Xiang Wang, Deli Zhao, and Jingren Zhou. I2vgen-xl: High-quality image-to-video synthesis via cascaded diffusion models, 2023. 2
- [62] Yuechen Zhang, Jinbo Xing, Eric Lo, and Jiaya Jia. Real-world image variation by aligning diffusion inversion chain. *arXiv preprint arXiv:2305.18729*, 2023. 2

- [63] Daquan Zhou, Weimin Wang, Hanshu Yan, Weiwei Lv, Yizhe Zhu, and Jiashi Feng. Magicvideo: Efficient video generation with latent diffusion models. *arXiv preprint arXiv:2211.11018*, 2022. [1](#), [2](#), [3](#)