



Full length article

AIGCBench: Comprehensive evaluation of image-to-video content generated by AI

Fanda Fan^{a,b}, Chunjie Luo^a, Wanling Gao^a, Jianfeng Zhan^{a,b,*}^a Research Center for Advanced Computer Systems, State Key Lab of Processors, Institute of Computing Technology, Chinese Academy of Sciences, China^b University of Chinese Academy of Sciences, China

ARTICLE INFO

Keywords:

Artificial intelligence generated content
Video generation
Image-to-video benchmark
Diffusion model
Multimodal AI

ABSTRACT

The burgeoning field of Artificial Intelligence Generated Content (AIGC) is witnessing rapid advancements, particularly in video generation. This paper introduces AIGCBench, a pioneering comprehensive and scalable benchmark designed to evaluate a variety of video generation tasks, with a primary focus on Image-to-Video (I2V) generation. AIGCBench tackles the limitations of existing benchmarks, which suffer from a lack of diverse datasets, by including a varied and open-domain image-text dataset that evaluates different state-of-the-art algorithms under equivalent conditions. We employ a novel text combiner and GPT-4 to create rich text prompts, which are then used to generate images via advanced Text-to-Image models. To establish a unified evaluation framework for video generation tasks, our benchmark includes 11 metrics spanning four dimensions to assess algorithm performance. These dimensions are control-video alignment, motion effects, temporal consistency, and video quality. These metrics are both reference video-based and video-free, ensuring a comprehensive evaluation strategy. The evaluation standard proposed correlates well with human judgment, providing insights into the strengths and weaknesses of current I2V algorithms. The findings from our extensive experiments aim to stimulate further research and development in the I2V field. AIGCBench represents a significant step toward creating standardized benchmarks for the broader AIGC landscape, proposing an adaptable and equitable framework for future assessments of video generation tasks. We have open-sourced the dataset and evaluation code on the project website: <https://www.benchcouncil.org/AIGCBench>.

1. Introduction

Artificial Intelligence Generated Content (AIGC) encompasses a wide array of applications that leverage AI technologies to automate the creation or editing of content across different media types, such as text, images, audio, and video. With the rapid advancement of diffusion models [1–5] and multimodal AI technologies [6], the AIGC field is experiencing considerable and rapid progress. The explosive growth of AIGC has made its evaluation and benchmarking an urgent task.

A representative application of AIGC is video generation [7–11]. Current video generation includes Text-to-Video (T2V), Image-to-Video (I2V), Video-to-Video (V2V), as well as a few other works that utilize additional information such as depth [10], pose [12], trajectory [13], and frequency [14] to generate videos. Among these, T2V and I2V are the two most mainstream tasks at present. Early video generation primarily used text prompts to generate videos and achieved good

results [7,8,15–19]. However, using text alone makes it difficult to depict the specific scenes that users want. Recently, I2V has ignited the AIGC community. The I2V task refers to the generation of a dynamic, moving video sequence based on a static input image and is usually accompanied by a text prompt.¹ Compared to T2V, I2V can better define the content of video generation, achieving excellent results in many scenarios such as film, e-commerce advertising, and micro-animation effects.

While benchmarks for the T2V task have seen notable progress [20–22], benchmarks for the I2V task have scarcely advanced. Previous efforts like Latent Flow Diffusion Models (LFDM) [23] and CATER-GEN [24] were tested under domain-specific video scenarios. VideoCrafter [25] and I2VGen-XL [26] only utilized visual comparisons for the I2V task. Seer [27] and Stable Video Diffusion (SVD) [28] employed video-text datasets and utilized a few metrics that require reference videos. Existing I2V benchmarks suffer from (1) a lack of

* Corresponding author at: Research Center for Advanced Computer Systems, State Key Lab of Processors, Institute of Computing Technology, Chinese Academy of Sciences, China.

E-mail addresses: fanfanda@ict.ac.cn (F. Fan), zhanjianfeng@ict.ac.cn (J. Zhan).

¹ However, the community often refers to it as Image-to-Video, rather than Text-Image-to-Video.

² Open-domain images refer to images that cover a wide variety of subjects or topics without specific restrictions on the content or category.

³ Here, equivalent conditions refer to using the same evaluation dataset and assessment dimensions for all video generation tasks.

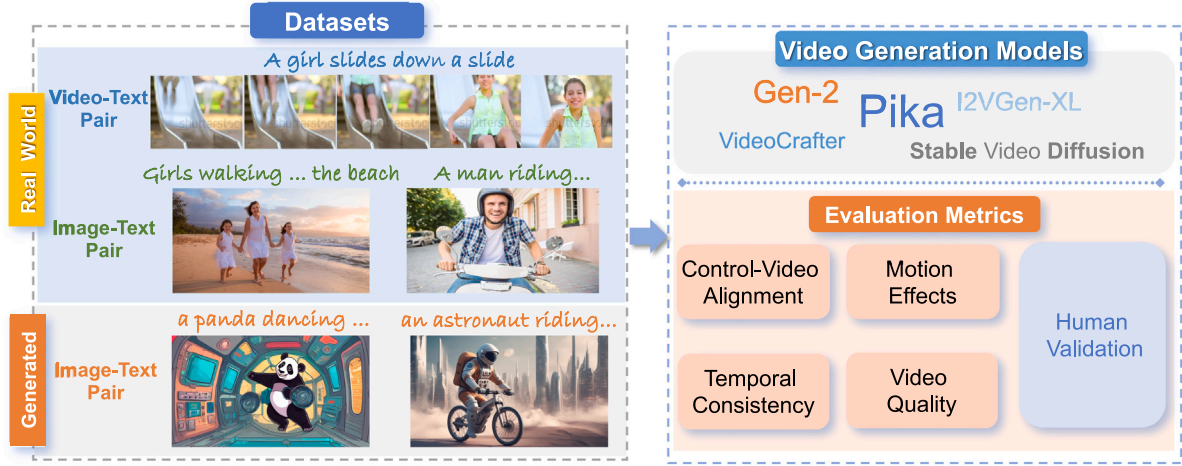


Fig. 1. Illustration of our AIGCBench. Our AIGCBench is divided into three modules: the evaluation dataset, the evaluation metrics, and the video generation models to be assessed. Our benchmark encompasses two types of datasets: video-text and image-text datasets. To construct a more comprehensive evaluation dataset, we expand the image-text dataset by our generation pipeline. Additionally, for a thorough evaluation of video generation models, we introduce a set of evaluation metrics comprising 11 metrics across four dimensions. These metrics include both reference video-based and reference video-free metrics, making full use of the benchmark we propose. We also adopted human validation to confirm the rationality of the evaluation standards we proposed.

diverse, open-domain images² with various subjects and styles to test the efficacy of different state-of-the-art algorithms; (2) an absence of a unified consensus on which evaluation metrics should be used to assess the final generated results. From the perspective of [29], these two shortcomings hinder the capability of capturing stakeholders' concerns and interests, while also failing to construct equivalent evaluation conditions.³

To address this gap, we present AIGCBench, a unified benchmark for video generation tasks. AIGCBench aims to encapsulate all mainstream video generation tasks, such as T2V, I2V, V2V, and the synthesis of video from additional modalities like depth, pose, trajectory, and frequency. We present an overview of AIGCBench in Fig. 1. Our AIGCBench is divided into three modules: the evaluation dataset, the evaluation metrics, and the video generation models to be assessed. Considering the high relevance and interconnectivity⁴ of video generation tasks, our AIGCBench can enable the comparison of different algorithms under equivalent evaluation conditions. This allows for an analysis of the strengths and weaknesses of different state-of-the-art video generation algorithms, thereby aiding progress in the field of video generation. In the first version of our AIGCBench, we address the current lack of a reasonable benchmark for I2V tasks by providing a thorough evaluation for them. In subsequent versions, we plan to include more video generation tasks and place them under equivalent evaluation conditions for a fair comparison.

Recognizing the limitations of existing benchmarks, AIGCBench is engineered to meet the diverse demands of users looking to animate a broad array of static images. Where previous benchmarks have fallen short, not fully accommodating the expansive range of images users might choose to animate – such as a blue dragon skateboarding in Times Square – AIGCBench rises to the challenge. We address this by deploying a text combiner to generate a rich assortment of text prompts that span a multitude of subjects, behaviors, backgrounds, and artistic styles. Further refining the creative process, we employ the advanced capabilities of GPT-4 [30] to enhance the text prompts, rendering them more vivid and intricate. These detailed prompts then guide the generation of images through state-of-the-art Text-to-Image diffusion models.

⁴ The interconnectivity arises because some algorithms have the capability to perform multiple types of video generation tasks.

By judiciously blending video-text and image-text datasets, along with our generated image-text pairs, AIGCBench ensures a robust and comprehensive evaluation of an array of I2V algorithms, thus addressing the first major shortcoming identified in existing benchmarks.

To establish a comprehensive and standardized set of evaluation metrics for video generation tasks that cater to mainstream tasks such as T2V and I2V, our AIGCBench evaluates four critical dimensions: control-video alignment, motion effects, temporal consistency, and video quality, thereby capturing every aspect of video generation. This integrated framework combines metrics that are both reference video-based and video-free metrics, enhancing the benchmark's rigor without exclusively relying on video-text datasets or image-text datasets alone. We strengthen this approach by incorporating image-text datasets into our evaluations, which allows us to assess content beyond the scope of existing video-text datasets and add reference video-free metrics for assessment. Considering the complexity and diversity of tasks, we believe that the evaluation metrics should cover at least these four aspects. For each aspect, we aim to use both reference video-based and video-free metrics. After satisfying these categorizations, the benefits of increasing the number of metrics become marginal, while it is insufficient without covering these aspects. The experimental results demonstrate that our evaluation standard correlates well with human ratings, confirming its effectiveness. Following a thorough evaluation, we present the strengths and weaknesses of each model, alongside several insightful findings, in hopes of spurring discussions that advance the I2V field.

Our contributions are as follows:

1. We introduce AIGCBench, a benchmark for comprehensive evaluation of diverse video generation tasks, with an initial focus on Image-to-Video (I2V) generation and a commitment to placing these models under equivalent evaluation conditions for fair comparison.
2. We extend our image-text dataset using a text combiner and GPT-4, complemented by state-of-the-art Text-to-Image models to generate high-quality images, enabling a deeper evaluation of I2V algorithm performance;
3. We evaluate I2V algorithms comprehensively using both reference video-based and video-free metrics across four aspects and verify the validity of our proposed evaluation standard with human judgment;

Table 1

Compare the features of our AIGCBench with other I2V benchmarks. ✗ and ✓ indicate whether the benchmark includes the features listed in the respective columns. Video-based metrics, which use reference videos, contrast with video-free metrics that do not. Considering the difficulty of the evaluation, we are not counting the sample numbers for domain-specific benchmarks [23,24,27].

Benchmark	Open-Domain	Video-Text Pairs	Image-Text Pairs	Generated Dataset	#Samples	Metric Types	# Metrics
LFDM Eval [23]	✗	✓	✗	✗	-	Video-based	3
CATER-GEN [24]	✗	✓	✓	✓	-	Video-based & Video-free	7
Seer Eval [27]	✗	✓	✗	✗	-	Video-based	2
VideoCrafter Eval [25]	✓	✓	✓	✗	-	-	-
I2VGen-XL Eval [26]	✓	✓	✓	✗	-	-	-
SVD Eval [28]	✓	✓	✓	✗	900	Video-based	5
AnimateBench [31]	✓	✗	✗	✓	105	Video-free	2
AIGCBench (Ours)	✓	✓	✓	✓	3928	Video-based & Video-free	11

4. We offer several insightful findings to aid the better development of the I2V community.

2. Background and related work

Current video generation primarily encompasses two major tasks: Text-to-Video (T2V) and Image-to-Video (I2V). Given the high relevance of T2V tasks to I2V tasks, we discuss video generation models, with a particular focus on I2V models, in Section 2.1. We will introduce related benchmarks for T2V in Section 2.2 and describe the existing benchmarks for I2V in Section 2.3.

2.1. Video generation models

Thanks to the development of diffusion models [1–5] and multimodal techniques [6], video generation algorithms are becoming increasingly sophisticated. Early video generation was primarily based on text-to-video approaches [7,8,10,15,17–19,32–36]. Most of the work is based on diffusion models [2–5,16], with some being transformer-based [15,33]. They all rely on extensive video-text or image-text datasets to train scalable models. However, considering that using only text can make it challenging to intuitively depict the video scenes users want to generate, image-to-video has started to gain popularity in the video generation community.

Seer [27] introduced an approach for I2V tasks that combines the conditional image latent with a noisy latent, utilizing causal attention within the temporal component of a 3D U-Net [37]. VideoComposer [38] concatenated image embedding with image style embedding to preserve the initial image information. Recently, VideoCrafter [25] encoded the image prompt through a lightweight image encoder and fed it into the cross-attention layer. Similarly, I2VGen-XL [26] not only merges the image latent with the noisy latent at the input layer but also employs a global encoder that extracts the image CLIP feature into the video latent diffusion model (VLDL). Stable video diffusion [28] is an extension of a pretrained image-based diffusion model [39]. It is trained through three stages: text-to-image pretraining, video pretraining, and high-quality video fine-tuning. Emu Video [40] identified critical design decisions, such as adjusted noise schedules for diffusion and multi-stage training, which enabled the generation of high-quality videos without requiring a deep cascade of models as in prior work. Beyond academic research, the video generation results from industry players like Pika [9] and Gen2 [10] are also quite impressive. All of these I2V algorithms are based on video diffusion models, and the majority leverage the parameter priors from image diffusion models to aid in the convergence of video models.

To evaluate state-of-the-art I2V models, we have reviewed three open-source works in this paper: VideoCrafter [25], I2VGen-XL [26], and Stable Video Diffusion [28], as well as two closed-source industry efforts, Pika [9] and Gen2 [10]. These currently represent the five

most influential works in the video generation community, and we will briefly introduce their experimental parameters in Section 5.1.

2.2. Benchmarks for text-to-video generation

The FETV benchmark [20] conducts a comprehensive manual evaluation of representative T2V models and reveals their strengths and weaknesses in handling a diverse range of text prompts from multiple perspectives. EvalCrafter [21] starts by creating a new set of prompts for T2V generation with the assistance of a large language model, ensuring that the prompts are representative of actual user queries. EvalCrafter’s benchmarks [21] are meticulously designed to evaluate generated videos from several critical dimensions: visual quality, content accuracy, motion dynamics, and the alignment between generated video content and the original text captions. VBench [22] has created 16 distinct evaluation dimensions, each with specialized prompts for precise assessment.

The task of T2V differs from I2V, as videos generated from the same text can vary widely, making it less suitable for evaluation metrics that require a reference video. For T2V tasks, the results generated by different models for the same text prompt can be quite dissimilar. However, for I2V tasks, since the image imposes certain constraints, the variation in results produced by different models is generally not as pronounced. This allows us to conduct a comprehensive evaluation of different Image-to-Video (I2V) algorithms on video-text datasets using evaluation metrics that are based on reference videos. Our AIGCBench draws on these T2V benchmarks but differs from them in several respects: (1). We need to collect or construct images for the I2V model’s input, which requires considering the comprehensiveness of both the text prompt set and the image set. (2). Although our evaluations are similar to those of the T2V task in terms of the dimensions assessed, we need to employ new evaluation standards due to the differences between T2V and I2V tasks.

2.3. Benchmarks for image-to-video generation

Domain-specific I2V benchmark. LFDM Eval [23] is evaluated on facial expression and human action datasets, employing just a few evaluation metrics to gauge the quality of video generation. The CATER-GEN [24] benchmark uses predefined 3D objects and specific initial images for testing the quality of videos that depict the motion of 3D objects. Nonetheless, neither LFDM Eval [23] nor the CATER-GEN [24] benchmark is appropriate for evaluating video generation in open-domain scenarios.

Open-domain I2V benchmark. The open-domain I2V benchmark is currently based on two main types of evaluation data: video-text and image-text datasets. Seer [27] and SVD [28] have utilized video-text datasets and employed a limited number of metrics that require reference videos for evaluation. VideoCrafter [25] and I2VGen-XL [26]

have used image-text datasets and relied solely on visual comparisons. Very recently, AnimateBench [31] was released for the purpose of evaluating I2V tasks. They also generated images using Text-to-Image models. However, they were limited by a small number of text prompts and a limited collection of images. At the same time, there is a lack of comprehensive evaluation metrics. Both are constrained by limited evaluation datasets and an incomplete set of assessment metrics. This leads to the evaluation datasets not being representative of all stakeholders' concerns and interests, and there is also a lack of a unified and comprehensive consensus on evaluation. In this paper, we expand the image-text dataset using state-of-the-art Text-to-Image models. To ensure the complexity of the generated text prompts, we generate prompts through the combinatorial traversal of four metatypes and enhance them with the capabilities of large language models. We compare our AIGCBench with other I2V benchmarks in Table 1.

Generating image-text dataset. While most benchmarks gather datasets from the real world, CATER-GEN [24] constructs datasets using a limited set of text prompts for specific object movement scenarios. Very recently, AnimateBench [31] utilized a limited number of manually designed text prompts and also employed Text-to-Image models to generate images. However, this approach is constrained by the simplicity of the text combinations and the limited diversity of the images. Our generation pipeline uses a text combiner to randomly generate text prompts and incorporates GPT-4 [30] to enrich the content. Simultaneously, we filter the generated results to select high-quality image-text pairs.

3. AIGCBench: Establishing the image-to-video generation benchmark

The framework of our AIGCBench is shown in Fig. 1. Our AIGCBench framework comprises three components: the evaluation dataset, the video generation models to be assessed, and the evaluation metrics. To construct a comprehensive benchmark, we evaluate I2V models using two types of datasets: video-text and image-text. For the image-text dataset, we utilize evaluation metrics that do not require reference videos. In this section, we will introduce how we collected the evaluation datasets, in Section 4 we present the evaluation criteria we have established, and in Section 5.1 we provide a brief introduction to the video generation models to be evaluated.

3.1. Collect dataset from real-world

Video-text pairs. The WebVid-10M [41] dataset is a substantial collection specifically designed to aid in the development and training of AI models for video understanding tasks. It consists of approximately 10 million video-text pairs, making it one of the larger datasets available for this type of research. Considering that video generation is time-consuming, we have sampled 1000 videos from the validation set of the WebVid10M [41] dataset based on subtype for evaluation purposes.

Image-text pairs. The LAION-5B [42] dataset is a large-scale, open dataset consisting of around 5.85 billion image-text pairs. It was created to facilitate research in computer vision and machine learning, specifically in areas such as multi-modal language-vision models, Text-to-Image generation, and more (e.g. CLIP [6], DALL-E [43]). LAION-Aesthetics is a subset from LAION-5B [42] with high visual quality. We randomly sampled 925 image-text pairs from the LAION-Aesthetics dataset to serve as a reference for video-free evaluation metrics.

3.2. Generated image-text pairs

Using only real-world datasets is insufficient. Users often input images and text generated by designers or T2I (Text-to-Image) models to create videos. This includes certain image-text pairs that cannot be sampled in the real world. To bridge this gap, we propose a T2I generation pipeline. As shown in Fig. 2, we provide an overview of our generation pipeline above and present some generated cases below.

3.2.1. Text combiner

To generate as diverse text prompts as possible, we construct text templates based on four types: subject, behavior, background, and image style. We then generate a list of 3000 text prompts randomly by following the template: **subject + behavior + background**, in the **image style** style. We have listed some examples:

1. Subject: *a dragon, a knight, an alien, a robot, a panda, a nymph;*
2. Behavior: *riding a bike, fight a monster, searching for a treasure, dancing, solving a puzzle;*
3. Background: *in a forest, in a futuristic city, in a space station, in an old western town at high noon;*
4. Image style: *oil painting, water color, cartoon, realistic, Van Gogh, Picasso.*

We have compiled our text corpus from high-frequency words often entered by users in the T2I community of Civit AI [44], along with some potentially valuable text prompts. Considering the flexibility of our generation pipeline, our benchmark is scalable. Subsequently, we can update and iterate on the versions of our text corpus.

3.2.2. Optimizing text prompts

Although utilizing text templates with various text corpora can generate reasonable images, it might lead to poor diversity in the generated images, which is not conducive to evaluating I2V tasks. We leverage the capabilities of the GPT-4 model [30], using the prompt “*make the content more vivid and rich*” to optimize the texts generated from templates.

3.2.3. Generate images and filter

To generate high-quality images based on the generated texts, we have employed the best Text-to-Image (T2I) model available to date — the Stable Diffusion model [39]. The Stable Diffusion model [39] is particularly notable for its ability to create high-quality and coherent images that closely match the style and content described by the input text prompts. We utilized the latest xl-base T2I model released by their community. Considering that the I2V model is primarily trained with an aspect ratio of 16:9, we used a height of 720 and a width of 1280 to generate images.

In order to select high-quality image-text pairs, we filtered out the top 2003 high-quality image-text pairs based on the automatic metrics from the T2I-CompBench [45]. Some examples generated by our pipeline can be seen in the lower half of Fig. 2.

4. Evaluation metrics

Our evaluation dataset includes both video-text and image-text datasets. To conduct a comprehensive evaluation, we employ two types of assessment metrics: one that requires reference videos and another that does not. In addition, we considered previous Text-to-Video benchmarks [20–22] and have integrated to propose an evaluation standard suitable for the Image-to-Video (I2V) task, covering both types of dataset. We assess the performance of different I2V models from four aspects: control-video alignment,⁵ motion effects, temporal consistency, and overall video quality⁶. Considering that videos generated by different algorithms have varying numbers of frames, for a standardized evaluation, we adopt the approach of extracting the first 16 frames, unless otherwise specified.

⁵ “control” refers to the input signals from the user, such as text, images, and other forms of control signals.

⁶ The code is available at <https://github.com/BenchCouncil/AIGCBench>.

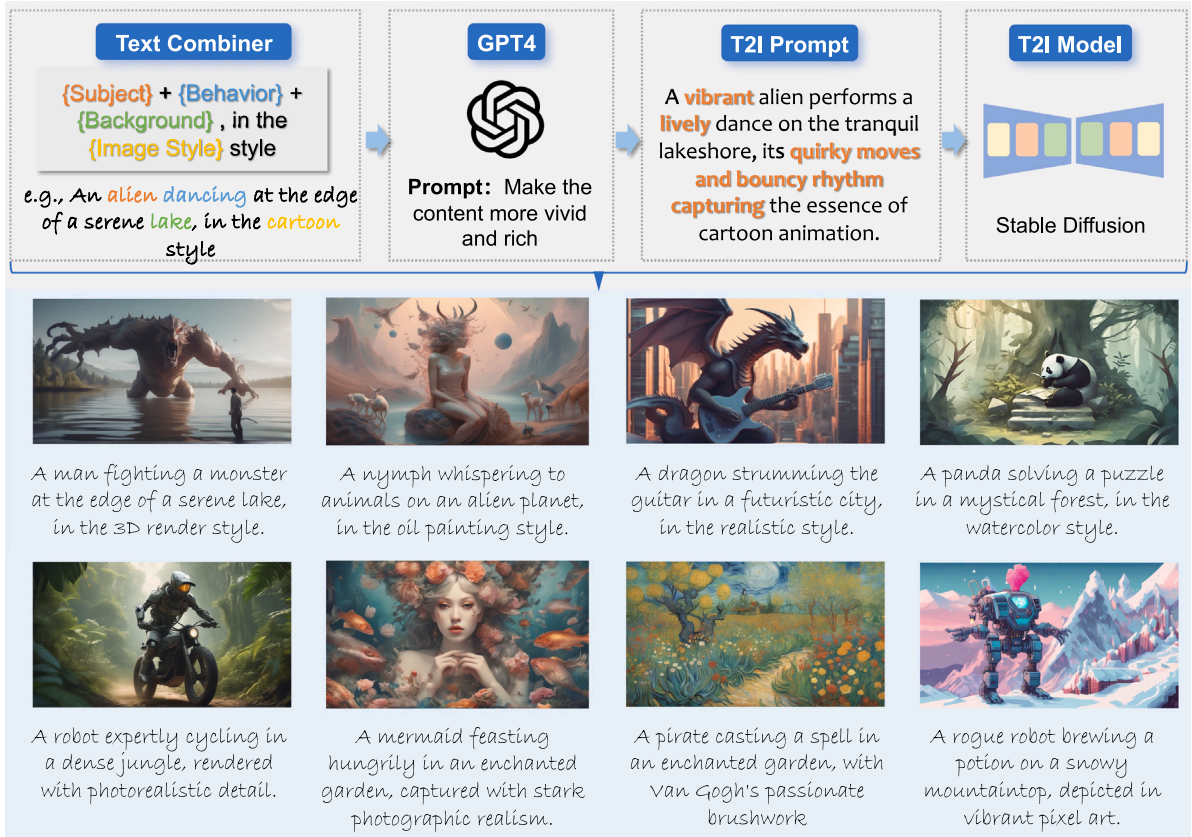


Fig. 2. Image-text dataset generation pipeline and results. Above: An overview of our T2I generation pipeline is presented. Below: Eight generated cases are showcased, with the original text produced by the text combiner displayed beneath each image.

4.1. Control-video alignment

The control-video alignment measures the degree of alignment between the user's input control signals, such as text and images, and the generated video. Considering that current video generation tasks primarily involve two types of inputs—a starting image and a text prompt—we introduce two evaluation metrics in the first version of our benchmark: **image fidelity** and **text-video alignment**. The image fidelity metric evaluates how similar the generated video frames are to the image input into the I2V model, especially the first frame. To assess fidelity, for the first frame of the generated video, we use metrics such as Mean Squared Error (MSE) and Structural Similarity Index Measure (SSIM) [46] to calculate the degree of preservation of the first frame. For all frames of the video, we calculate the similarity of CLIP (Contrastive Language-Image Pre-training) [6] embeddings between the input image and each frame of the generated video. We use MSE (First), SSIM (First), and Image-GenVideo CLIP to represent these three evaluation metrics, respectively.

Considering that the I2V models we evaluate also take text as input, we need to assess whether the generated videos are relevant to the input text. For the generated videos, we use CLIP [6] to calculate the similarity between the input text and the generated video results. We assume that the videos in the video-text dataset are consistent with the textual descriptions. For the video-text dataset, we use the keyframes from the reference videos and the generated videos to compute the CLIP [6] similarity. Considering that the text typically describes high-level semantics and that the generated videos may not correspond perfectly with the original videos, we uniformly sample four keyframes for comparison. We use GenVideo-Text Clip and GenVideo-RefVideo CLIP (Keyframes) to represent these two evaluation metrics, respectively.

4.2. Motion effects

Motion effects primarily evaluate whether the amplitude of the motion in the generated video is significant and whether the movements are reasonable. As for the amplitude of the motion, we follow the [21, 22] and use a pretrained optical flow estimation method, RAFT [47], to calculate the flow score between adjacent frames of the generated video, with the final average value representing the magnitude of the motion effects. We use the square average of the predicted values from adjacent frames to represent the motion dynamics of the video, with higher values indicating stronger motion effects. Considering that there are some bad cases in video generation, we set a threshold where the square average value must be less than 10 to filter out these bad cases. For the video-text dataset, we have real videos corresponding to the text. We measure the reasonableness of the generated motion effects by calculating the similarity between each frame of the generated video and each frame of the reference video, and then taking the average. For robustness, we use the image CLIP [6] metric to calculate the similarity between frames. We use Flow-Square-Mean and GenVideo-RefVideo CLIP (Corresponding frames) to represent these two evaluation metrics, respectively.

4.3. Temporal consistency

Temporal consistency measures whether the generated video frames are consistent and coherent with each other. We calculate the image CLIP [6] similarity between every two adjacent frames in the generated video and take the average as an indicator of the temporal consistency of the generated video. We use GenVideo Clip (Adjacent frames) to represent this evaluation metric. In addition, we also use GenVideo-RefVideo (Corresponding frames) from Section 4.2 to represent temporal consistency.

Table 2

Quantitative analysis for different Image-to-Video algorithms. An upward arrow indicates that higher values are better, while a downward arrow means lower values are preferable.

Dimensions	Metrics	VideoCrafter [25]	I2VGen-XL [26]	SVD [28]	Pika [9]	Gen2 [10]
Control-video alignment	MSE (First) ↓	3929.65	4491.90	640.75	155.30	235.53
	SSIM (First) ↑	0.300	0.354	0.612	0.800	0.803
	Image-GenVideo Clip ↑	0.830	0.832	0.919	0.930	0.939
	GenVideo-Text Clip ↑	0.23	0.24	–	0.271	0.270
	GenVideo-RefVideo CLiP (Keyframes) ↑	0.763	0.764	–	0.824	0.820
Motion effects	Flow-Square-Mean	1.24	1.80	2.52	0.281	1.18
	GenVideo-RefVideo CLiP (Corresponding frames) ↑	0.764	0.764	0.796	0.823	0.818
Temporal consistency	GenVideo Clip (Adjacent frames) ↑	0.980	0.971	0.974	0.996	0.995
	GenVideo-RefVideo CLiP (Corresponding frames) ↑	0.764	0.764	0.796	0.823	0.818
Video quality	Frame Count ↑	16	32	25	72	96
	DOVER ↑	0.518	0.510	0.623	0.715	0.775
	GenVideo-RefVideo SSIM ↑	0.367	0.304	0.507	0.560	0.504

4.4. Video quality

Video quality is a relatively subjective dimension, measuring the overall quality of video production. We first use the number of frames generated by videos to gauge the ability of different algorithms to generate long videos. We utilize disentangled objective video quality evaluator (DOVER) [48], a no-reference video quality assessment metric. DOVER [48] comprehensively rates videos from both aesthetic and technical perspectives, using the collected DIVIDE-3k dataset. Experimental results show that the DOVER [48] metric highly correlates with human opinions in both aesthetic and technical perspectives. For the DOVER evaluation metric, we calculate it using all frames generated by their respective algorithms. For the video-text dataset, since we have reference videos available, we measure the spatial structural similarity of the generated videos to the reference videos by calculating the SSIM (Structural Similarity Index Measure) between the corresponding frames of the generated and reference videos. We denote this evaluation metric as GenVideo-RefVideo SSIM.

5. Experiments

5.1. Evaluated models

5.1.1. Open-source project

VideoCrafter. VideoCrafter [25] is an open-source video generation and editing toolbox for crafting video content. It supports the generation of videos from images. We use a guidance scale of 12 and ddim steps of 25. For videos with an aspect ratio of 1, we employ a resolution of 512 * 512, while for videos with an aspect ratio of 0.5625, we use a resolution of 512 * 320, and then uniformly resize to align with the resolutions used by other methods.

I2Vgen-XL. I2VGen-XL [26] is an open-source video synthesis codebase developed by Tongyi Lab at Alibaba Group, which features state-of-the-art video generative models. We use a guide scale of 9 and infer with fp16 precision.

Stable video diffusion. Stable Video Diffusion (SVD) [28] is an expansion of the model based on Image Stable Diffusion [39]. We use the 25-frame version of Stable Video Diffusion. It is worth noting that the current model does not support text input temporarily, hence we did not calculate the text-video alignment for this model.

5.1.2. Closed-source project

Pika. Pika [9] is a technology company revolutionizing video creation by making it effortless and accessible for everyone. In just six months, Pika has built a community of half a million users producing millions of videos per week. The company recently launched Pika 1.0, a significant upgrade featuring a new AI model that supports various video styles,

including 3D animation, anime, cartoons, and cinematic, coupled with an improved web experience. Considering that Pika [9] does not have open-source code, we manually tested 60 cases on the Discord platform (30 from the WebVid dataset and 30 from our own generated dataset). We used the default parameters of motion set to 1 and the guidance scale set to 12.

Gen2. Gen2 [10] is a multimodal AI system that can generate novel videos with text, images, or video clips. We used the default motion setting of 5 from the demo and did not employ the camera movement parameter to generate videos.

5.2. Comprehensive results analysis

Table 2 presents the evaluation of five state-of-the-art (SOTA) I2V algorithms across five dimensions: image fidelity, motion effects, text-video alignment, temporal consistency, and video quality. We present the qualitative results of different I2V algorithms in Fig. 3. We find that VideoCrafter and I2VGen-xl struggle to preserve the original image. I2VGen-xl maintains relatively good semantics, but the spatial structure of the initial image is mostly not preserved. VideoCrafter can approximate the spatial structure of the initial image to some extent, but the preservation of details is generally mediocre. SVD, Pika, and Gen2 preserve the original image quite well, with Gen2 achieving the best preservation effect. As for the aspect of Text-video alignment, Gen2 and Pika are nearly on par with each other and both outperform the open-source algorithms. However, existing algorithms and evaluation metrics do not effectively capture fine-grained textual changes. In terms of motion effects, VideoCrafter tends to remain static. I2VGen-xl and SVD lean towards camera movement rather than subject motion, which is why they score high on the flow-square-mean but obtain low GenVideo-RefVideo Clip scores. Pika tends to favor both local and subject movement, thus achieving high GenVideo-RefVideo Clip scores and low flow-square-mean scores. Gen2, on the other hand, favors movement in both the foreground and background, but the background movement is not as pronounced as with SVD.

In the aspect of temporal consistency, VideoCrafter, due to its poorer motion effects, does not perform poorly in terms of temporal consistency. Considering that SVD has stronger motion effects and still maintains good temporal consistency, it has achieved the best performance among open-source I2V algorithms. Similarly, Pika, because of its tendency for local movement, has achieved the highest score in overall temporal consistency. As for video quality, Gen2 is capable of generating the longest videos of up to 96 frames, with the highest levels of aesthetics and clarity. Pika, due to its tendency for local movement, has achieved the highest similarity in the GenVideo-RefVideo SSIM metric. SVD benefits from the priors of the image stable diffusion model, resulting in videos that reach the best performance among open-source I2V algorithms. In summary, the two closed-source projects, Pika

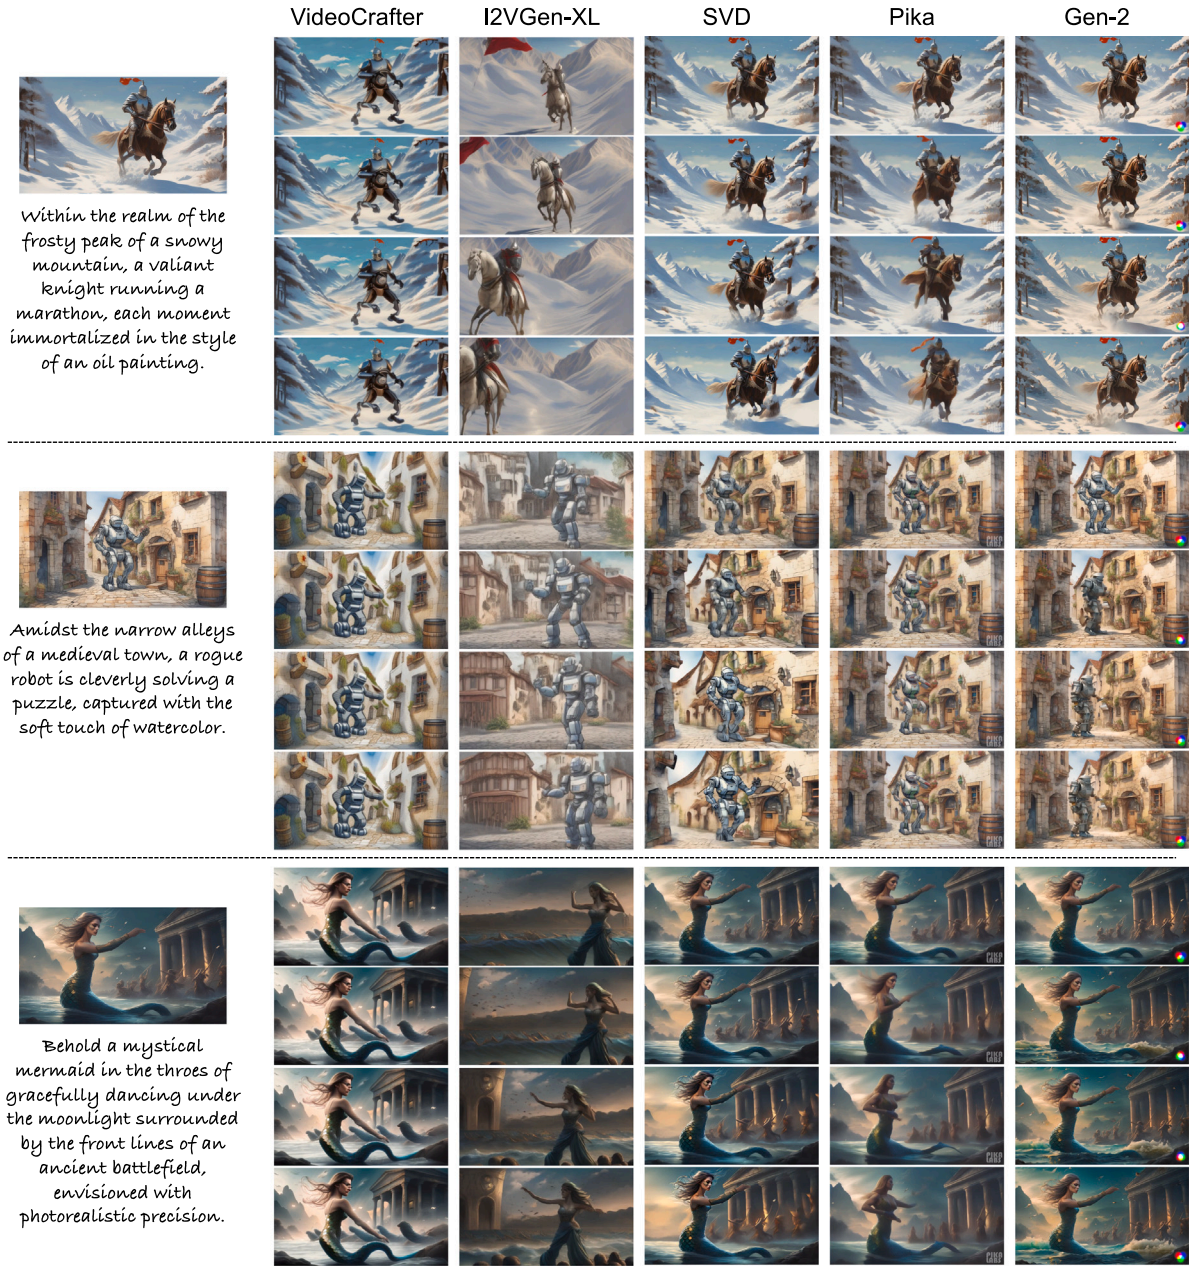


Fig. 3. We present three I2V cases utilizing five state-of-the-art algorithms, among which VideoCrafter, I2VGen-XL, and SVD are open-source research, while Pika and Gen2 are closed-source project. For additional videos, please refer to our supplementary materials.

and Gen2, achieved the most optimal generation effects, capable of producing long videos. Pika excels in generating local motion, while Gen2 tends to prefer global motion. SVD achieved the best results among the open-source options, demonstrating outcomes that were close to those of the two closed-source projects.

5.3. User study

To validate whether the proposed evaluation standards are aligned with human preference, we randomly sampled 30 generated results from each of the five methods and tallied the best algorithm outcomes in each of the four dimensions (Image Fidelity, Motion Effects, Temporal Consistency, Video Quality) through human voting. We have tallied the votes of a total of 42 individuals, with the specific results presented in Fig. 4. We discovered that Gen2's performance is on par with Pika, both achieving optimal results. Pika excelled in temporal consistency and motion effects, while Gen2 came out on top in terms of image

fidelity and video quality. SVD showed a balanced performance across all areas, securing the best results among the open-source options. We found that the users' votes are relatively consistent with the results evaluated by our assessment criteria.

5.4. Findings and discussions

Despite the notable achievements of I2V and the rapid updates of new algorithms, there is still significant room for improvement in existing solutions. Utilizing AIGCBench, we have conducted a detailed survey and evaluation of the five most advanced I2V algorithms from both academia and industry. The comprehensive analysis facilitated by AIGCBench has led us to make the following discoveries:

Lacking fine-grained control. AIGCBench's testing methodology has revealed that current I2V tasks often fall short in allowing users to generate content with precise textual descriptions. The benchmark's

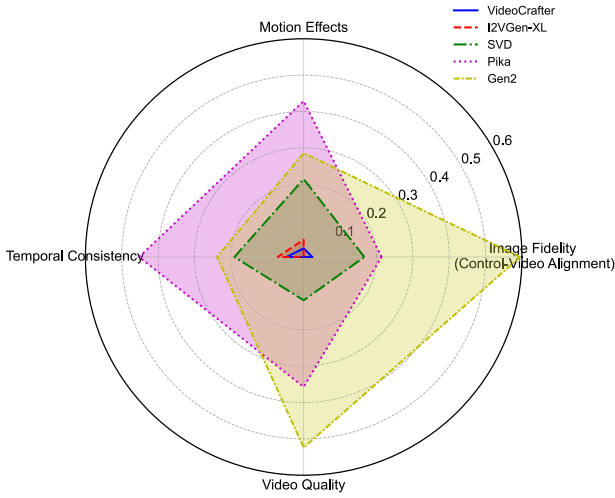


Fig. 4. We tallied the votes of 42 individuals, evaluating five state-of-the-art I2V algorithms from four aspects. The numerical values in the radar chart represent the proportion of users who voted for each algorithm as being the best performer in that aspect.

diverse datasets and nuanced evaluation metrics have shown that while solutions based on CLIP [6] and large language models [49] are a step in the right direction, they are not fully capturing the fine-grained details that yield a high degree of control over video content. AIGCBench’s detailed feedback on these algorithms has underscored the need for a model that is specifically trained for video contexts to improve text-video alignment and to provide users with that control. Our findings suggest that integrating AIGCBench’s evaluation criteria into the development process could lead to algorithms that better align with human preferences.

Long video generation. The current I2V algorithms can generate up to 96 frames in a single inference, which is far from satisfying users’ needs for longer video production. Considering that video scenes typically have a frame rate of 24 fps, the basic generation capability of mainstream algorithms is around 3 s. There are mainly two approaches to address this limitation. One is to use multiple inferences, where most adopt a coarse-to-fine generation pipeline—first generating diluted keyframes, then densely producing all frames. The challenge of this method lies in maintaining temporal consistency across multiple inferences. The other approach is to use multi-GPU training and inference with a single model, which currently struggles to guarantee satisfactory results. How to generate longer videos should be an urgent issue for the AI-generated content (AIGC) community to address next.

Inference speed. Currently, the speed of video generation is relatively slow. For a 3-second video, mainstream algorithms generally require about 1 min on a V100 graphics card. Considering that video generation scenarios are based on diffusion models [1–5], there are currently two main routes for speeding up the process. One is to reduce the dimensionality of the video in the latent space. For example, Stable Diffusion [28] maps the video into a latent space, roughly decreasing the size of the video by about 8 times, with only a minimal loss of video quality. The other is to improve the inference speed of the diffusion model, which is also a hot research topic in the AIGC community.

In light of these findings from AIGCBench, it becomes evident that the benchmark not only aids in the identification of current shortcomings but also offers a structure for addressing them. AIGCBench’s comprehensive framework for evaluation, encompassing a rich variety of datasets and multi-dimensional metrics, provides a roadmap for advancing the state of I2V algorithms. By exposing algorithmic weaknesses and offering a standardized platform for comparison, AIGCBench guides researchers towards developing solutions that overcome the

challenges of fine-grained control, video length, and inference speed. The insights gained from AIGCBench evaluations are instrumental in pushing the boundaries of what is possible in the field of I2V generation.

6. Conclusion

In this work, we have introduced AIGCBench, a comprehensive and scalable benchmark tailored for the evaluation of Image-to-Video (I2V) generation tasks. AIGCBench provides a much-needed framework to assess the performance of various state-of-the-art I2V algorithms under equivalent evaluation conditions. Our benchmark stands out by incorporating a diverse set of real-world video-text and image-text datasets, as well as a novel dataset produced through our proprietary generation pipeline. We have also proposed a novel set of evaluation metrics that span across four critical dimensions: control-video alignment, motion effects, temporal consistency, and video quality. These metrics have been validated against human judgment to ensure their alignment with human preferences. Our extensive evaluation of leading I2V models has not only highlighted their strengths and weaknesses but also unearthed significant insights that will guide the future development of the I2V domain.

AIGCBench marks a foundational step in benchmarking for AIGC, pushing the frontier of I2V technology evaluation. By offering a scalable and precise assessment methodology, we set the stage for continuous enhancements and innovations in this rapidly evolving research field. As we progress, we plan to expand AIGCBench to encompass a broader range of video generation tasks, creating a unified and extensive benchmark that reflects the multifaceted nature of AIGC.

7. Limitations and future work

Due to the slow inference speed of video generation by I2V models and the fact that some works are not open-sourced (e.g., Pika [9], Gen2 [10]), our benchmark only evaluated 3928 test cases. Considering the complexity of video generation tasks, we believe this number is insufficient. Furthermore, given the lack of fine-grained video recognition models currently available, our evaluation system is unable to accurately judge whether the direction of object movement in the generated videos matches the text description. For instance, whether water flows from left to right or from right to left, we are currently unable to determine through automated evaluation metrics if the direction of the water flow in the generated video is consistent with the textual description.

Moving forward, we will integrate tasks related to T2V and new video generation tasks into a large-scale video generation benchmark. Additionally, to address the issues mentioned above, we may train a fine-grained video representation model aligned with text, which will be utilized for fine-grained alignment of video and text scenes.

CRedit authorship contribution statement

Fanda Fan: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Chunjie Luo:** Writing – review & editing, Writing – original draft, Resources, Investigation, Conceptualization. **Wanling Gao:** Writing – review & editing, Writing – original draft, Methodology, Investigation. **Jianfeng Zhan:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Methodology, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

Fanda Fan, Chunjie Luo, Wanling Gao and Jianfeng Zhan are faculty at Institute of Computing Technology, Chinese Academy of Sciences University of Chinese Academy of Sciences.

Acknowledgments

We want to thank the reviewers and editors for their constructive comments and suggestions. This research is partly supported by the Strategic Priority Research Program of the Chinese Academy of Sciences, Grant No. XDA0320000 and XDA0320300. I would like to extend my heartfelt thanks to Professor Lei Wang for his insightful discussions and valuable revisions to this manuscript. I am also grateful to Mengya He for her contributions to the discussions of this paper. Furthermore, I wish to acknowledge Litong Gong, Weijie Li, Yiran Zhu, and Biao Wang from Alibaba Company for their support in the experimental aspects of this research.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.tbench.2024.100152>.

References

- [1] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, S. Ganguli, Deep unsupervised learning using nonequilibrium thermodynamics, in: *International Conference on Machine Learning*, PMLR, 2015, pp. 2256–2265.
- [2] Y. Song, S. Ermon, Generative modeling by estimating gradients of the data distribution, in: *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [3] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, in: *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 6840–6851.
- [4] A.Q. Nichol, P. Dhariwal, Improved denoising diffusion probabilistic models, in: *International Conference on Machine Learning*, PMLR, 2021, pp. 8162–8171.
- [5] P. Dhariwal, A. Nichol, Diffusion models beat gans on image synthesis, in: *Advances in Neural Information Processing Systems*, vol. 34, 2021, pp. 8780–8794.
- [6] A. Radford, J.W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al., Learning transferable visual models from natural language supervision, in: *International Conference on Machine Learning*, PMLR, 2021, pp. 8748–8763.
- [7] U. Singer, A. Polyak, T. Hayes, X. Yin, J. An, S. Zhang, Q. Hu, H. Yang, O. Ashual, O. Gafni, et al., Make-a-video: Text-to-video generation without text-video data, 2022, arXiv preprint [arXiv:2209.14792](https://arxiv.org/abs/2209.14792).
- [8] J. Ho, W. Chan, C. Saharia, J. Whang, R. Gao, A. Gritsenko, D.P. Kingma, B. Poole, M. Norouzi, D.J. Fleet, et al., Imagen video: High definition video generation with diffusion models, 2022, arXiv preprint [arXiv:2210.02303](https://arxiv.org/abs/2210.02303).
- [9] I. Pika, Pika lab discord server, 2023, <https://www.pika.art/>. (Accessed: 30 August 2023).
- [10] P. Esser, J. Chiu, P. Atighehchian, J. Granskog, A. Germanidis, Structure and content-guided video synthesis with diffusion models, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 7346–7356.
- [11] Q. Sun, Y. Cui, X. Zhang, F. Zhang, Q. Yu, Z. Luo, Y. Wang, Y. Rao, J. Liu, T. Huang, X. Wang, Generative multimodal models are in-context learners, 2023, arXiv:2312.13286.
- [12] J. Karras, A. Holynski, T.-C. Wang, I. Kemelmacher-Shlizerman, Dreampose: Fashion image-to-video synthesis via stable diffusion, 2023, arXiv preprint [arXiv:2304.06025](https://arxiv.org/abs/2304.06025).
- [13] S. Yin, C. Wu, J. Liang, J. Shi, H. Li, G. Ming, N. Duan, Dragnuwa: Fine-grained control in video generation by integrating text, image, and trajectory, 2023, arXiv preprint [arXiv:2308.08089](https://arxiv.org/abs/2308.08089).
- [14] Z. Li, R. Tucker, N. Snively, A. Holynski, Generative image dynamics, 2023, arXiv preprint [arXiv:2309.07906](https://arxiv.org/abs/2309.07906).
- [15] W. Hong, M. Ding, W. Zheng, X. Liu, J. Tang, Cogvideo: Large-scale pretraining for text-to-video generation via transformers, 2022, arXiv preprint [arXiv:2205.15868](https://arxiv.org/abs/2205.15868).
- [16] Y. He, T. Yang, Y. Zhang, Y. Shan, Q. Chen, Latent video diffusion models for high-fidelity video generation with arbitrary lengths, 2022, arXiv preprint [arXiv:2211.13221](https://arxiv.org/abs/2211.13221).
- [17] J.Z. Wu, Y. Ge, X. Wang, S.W. Lei, Y. Gu, Y. Shi, W. Hsu, Y. Shan, X. Qie, M.Z. Shou, Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 7623–7633.
- [18] Z. Luo, D. Chen, Y. Zhang, Y. Huang, L. Wang, Y. Shen, D. Zhao, J. Zhou, T. Tan, VideoFusion: Decomposed diffusion models for high-quality video generation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 10209–10218.
- [19] Y. Guo, C. Yang, A. Rao, Y. Wang, Y. Qiao, D. Lin, B. Dai, Animatediff: Animate your personalized text-to-image diffusion models without specific tuning, 2023, arXiv preprint [arXiv:2307.04725](https://arxiv.org/abs/2307.04725).
- [20] Y. Liu, L. Li, S. Ren, R. Gao, S. Li, S. Chen, X. Sun, L. Hou, FETV: A benchmark for fine-grained evaluation of open-domain text-to-video generation, 2023, arXiv preprint [arXiv:2311.01813](https://arxiv.org/abs/2311.01813).
- [21] Y. Liu, X. Cun, X. Liu, X. Wang, Y. Zhang, H. Chen, Y. Liu, T. Zeng, R. Chan, Y. Shan, EvalCrafter: Benchmarking and evaluating large video generation models, 2023, arXiv preprint [arXiv:2310.11440](https://arxiv.org/abs/2310.11440).
- [22] Z. Huang, Y. He, J. Yu, F. Zhang, C. Si, Y. Jiang, Y. Zhang, T. Wu, Q. Jin, N. Chanpaisit, et al., VBench: Comprehensive benchmark suite for video generative models, 2023, arXiv preprint [arXiv:2311.17982](https://arxiv.org/abs/2311.17982).
- [23] H. Ni, C. Shi, K. Li, S.X. Huang, M.R. Min, Conditional image-to-video generation with latent flow diffusion models, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 18444–18455.
- [24] Y. Hu, C. Luo, Z. Chen, A benchmark for controllable text-image-to-video generation, *IEEE Trans. Multimed.* (2023).
- [25] H. Chen, M. Xia, Y. He, Y. Zhang, X. Cun, S. Yang, J. Xing, Y. Liu, Q. Chen, X. Wang, et al., Videocrafter1: Open diffusion models for high-quality video generation, 2023, arXiv preprint [arXiv:2310.19512](https://arxiv.org/abs/2310.19512).
- [26] S. Zhang, J. Wang, Y. Zhang, K. Zhao, H. Yuan, Z. Qin, X. Wang, D. Zhao, J. Zhou, I2vgen-xl: High-quality image-to-video synthesis via cascaded diffusion models, 2023, arXiv preprint [arXiv:2311.04145](https://arxiv.org/abs/2311.04145).
- [27] X. Gu, C. Wen, J. Song, Y. Gao, Seer: Language instructed video prediction with latent diffusion models, 2023, arXiv preprint [arXiv:2303.14897](https://arxiv.org/abs/2303.14897).
- [28] A. Blattmann, T. Dockhorn, S. Kulal, D. Mendelevitch, M. Kilian, D. Lorenz, Y. Levi, Z. English, V. Voleti, A. Letts, et al., Stable video diffusion: Scaling latent video diffusion models to large datasets, 2023, arXiv preprint [arXiv:2311.15127](https://arxiv.org/abs/2311.15127).
- [29] J. Zhan, L. Wang, W. Gao, H. Li, Y. Huang, C. Wang, Y. Li, Z. Yang, G. Kang, C. Luo, H. Ye, S. Dai, Z. Zhang, Evaluatool: The Science and Engineering of Evaluation, Technical Report, Institute of Computing Technology Chinese Academy of Sciences, 2024.
- [30] OpenAI, GPT-4 technical report, 2023, arXiv:2303.08774.
- [31] Y. Zhang, Z. Xing, Y. Zeng, Y. Fang, K. Chen, PIA: Your personalized image animator via plug-and-play modules in text-to-image models, 2023, arXiv preprint [arXiv:2312.13964](https://arxiv.org/abs/2312.13964).
- [32] A. Blattmann, R. Rombach, H. Ling, T. Dockhorn, S.W. Kim, S. Fidler, K. 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*, 2023, pp. 22563–22575.
- [33] L. Yu, Y. Cheng, K. Sohn, J. Lezama, H. Zhang, H. Chang, A.G. Hauptmann, M.-H. Yang, Y. Hao, I. Essa, L. Jiang, MAGVIT: Masked generative video transformer, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [34] S. Ge, S. Nah, G. Liu, T. Poon, A. Tao, B. Catanzaro, D. Jacobs, J.-B. Huang, M.-Y. Liu, Y. Balaji, Preserve your own correlation: A noise prior for video diffusion models, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 22930–22941.
- [35] L. Khachatryan, A. Movsisyan, V. Tadevosyan, R. Henschel, Z. Wang, S. Navasardyan, H. Shi, Text2video-zero: Text-to-image diffusion models are zero-shot video generators, 2023, arXiv preprint [arXiv:2303.13439](https://arxiv.org/abs/2303.13439).
- [36] J. Wang, H. Yuan, D. Chen, Y. Zhang, X. Wang, S. Zhang, Modelscope text-to-video technical report, 2023, arXiv preprint [arXiv:2308.06571](https://arxiv.org/abs/2308.06571).
- [37] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, Springer, 2015, pp. 234–241.
- [38] X. Wang, H. Yuan, S. Zhang, D. Chen, J. Wang, Y. Zhang, Y. Shen, D. Zhao, J. Zhou, VideoComposer: Compositional video synthesis with motion controllability, 2023, arXiv preprint [arXiv:2306.02018](https://arxiv.org/abs/2306.02018).
- [39] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10684–10695.
- [40] R. Girdhar, M. Singh, A. Brown, Q. Duval, S. Azadi, S.S. Rambhatla, A. Shah, X. Yin, D. Parikh, I. Misra, Emu video: Factorizing text-to-video generation by explicit image conditioning, 2023, arXiv preprint [arXiv:2311.10709](https://arxiv.org/abs/2311.10709).
- [41] M. Bain, A. Nagrani, G. Varol, A. Zisserman, Frozen in time: A joint video and image encoder for end-to-end retrieval, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 1728–1738.
- [42] C. Schuhmann, R. Beaumont, R. Vencu, C. Gordon, R. Wightman, M. Cherti, T. Coombes, A. Katta, C. Mullis, M. Wortsman, P. Schramowski, S. Kundurthy, K. Crowson, L. Schmidt, R. Kaczmarczyk, J. Jitsev, LAION-5B: An open large-scale dataset for training next generation image-text models, 2022, arXiv:2210.08402.
- [43] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, I. Sutskever, Zero-shot text-to-image generation, in: *International Conference on Machine Learning*, PMLR, 2021, pp. 8821–8831.

- [44] I. Civit AI, CivitAI, 2022, <https://civitai.com/> (Accessed: (2022)).
- [45] K. Huang, K. Sun, E. Xie, Z. Li, X. Liu, T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation, 2023, arXiv preprint [arXiv:2307.06350](https://arxiv.org/abs/2307.06350).
- [46] Z. Wang, A.C. Bovik, H.R. Sheikh, E.P. Simoncelli, Image quality assessment: From error visibility to structural similarity, *IEEE Trans. Image Process.* 13 (4) (2004) 600–612.
- [47] Z. Teed, J. 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*, Springer, 2020, pp. 402–419.
- [48] H. Wu, E. Zhang, L. Liao, C. Chen, J. Hou, A. Wang, W. Sun, Q. Yan, W. Lin, Exploring video quality assessment on user generated contents from aesthetic and technical perspectives, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 20144–20154.
- [49] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P.J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, *J. Mach. Learn. Res.* 21 (1) (2020) 5485–5551.