

GenVidBench: A 6-Million Benchmark for AI-Generated Video Detection

Zhenliang Ni*, Qiangyu Yan*, Mouxiao Huang, Tianning Yuan, Yehui Tang,
Hailin Hu , Xinghao Chen , Yunhe Wang 

Huawei Noah's Ark Lab

{nizhenliang2, qiangyu.yan, hailin.hu, xinghao.chen, yunhe.wang}@huawei.com

Abstract

The rapid advancement of video generation models has made it increasingly challenging to distinguish AI-generated videos from real ones. This issue underscores the urgent need for effective AI-generated video detectors to prevent the dissemination of false information via such videos. However, the development of high-performance AI-generated video detectors is currently impeded by the lack of large-scale, high-quality datasets specifically designed for generative video detection. To this end, we introduce GenVidBench, a challenging AI-generated video detection dataset with several key advantages: 1) Large-scale video collection: The dataset contains 6.78 million videos and is currently the largest dataset for AI-generated video detection. 2) Cross-Source and Cross-Generator: The cross-source generation reduces the interference of video content on the detection. The cross-generator ensures diversity in video attributes between the training and test sets, preventing them from being overly similar. 3) State-of-the-Art Video Generators: The dataset includes videos from 11 state-of-the-art AI video generators, ensuring that it covers the latest advancements in the field of video generation. These generators ensure that the datasets are not only large in scale but also diverse, aiding in the development of generalized and effective detection models. Additionally, we present extensive experimental results with advanced video classification models. With GenVidBench, researchers can efficiently develop and evaluate AI-generated video detection models.

Code and Datasets — <https://genvidbench.github.io/>

Introduction

In recent years, video generation models like Sora have seen remarkable advancements (OpenAI 2024; Blattmann et al. 2023; Xia et al. 2024), leading to a significant enhancement in the quality of AI-generated videos. The line between realistically generated videos and real videos has become increasingly blurred, posing challenges such as the spread of misinformation, damage to personal and corporate reputations, and an escalation in cybersecurity threats (Bai, Lin, and Cao 2024; Ma et al. 2024). To address these risks, there is an urgent demand for the development of AI-generated




Figure 1: Overview of the proposed GenVidBench dataset. In addition to real/fake video labels, GenVidBench provides rich semantic annotations, including object categories, locations, and actions.

video detectors that can accurately identify and differentiate between real and fake videos. However, the development and evaluation of such detectors are still hindered by the lack of large-scale, challenging datasets. To this end, we present GenVidBench, a comprehensive and challenging dataset for the development of AI-generated video detectors.

GenVidBench is the first 6-million dataset for AI-generated video detection, and it features cross-source and cross-generator characteristics. It covers 11 types of videos generated by state-of-the-art video generators, such as Mora, MuseV, Sora, and Kling, ensuring that the generated videos are of exceptional quality. GenVidBench contains two real video sources: Vript and HD-VG-130M. To make the dataset challenging, we construct two sets of paired videos, each derived from the same text prompt or image. The first video pair contains Pika, VideoCraftV2, Modelscope, and T2V-zero, which are included in the training set. The second video pair, which contains HD-VG-130M, MuseV, SVD, Mora, CogVideo, is included in the test set. For the second video pair, we generate prompts and images based on the real video and use Image-to-Video (I2V) models and Text-to-Video (T2V) models to generate videos from the same source. Videos with the same generation source have more similar attributes, which makes it highly difficult to distinguish real videos from AI-generated videos. In GenVidBench, the two sets of paired videos are

* Equal contribution.

 Corresponding authors.

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

Dataset	Scale	Prompt/Image	Video Pairs	Semantic Label	Cross Source
GVD (Bai, Lin, and Cao 2024)	11k	×	×	×	×
GVF (Ma et al. 2024)	2.8k	✓	✓	✓	×
GenVideo (Chen et al. 2024a)	2.27M	×	×	×	×
GenVidDet (Ji et al. 2024)	2.66M	×	×	×	×
GenVidBench	6.78M	✓	✓	✓	✓

Table 1: An overview of fake video detection datasets. The proposed GenVidBench is the first dataset with a scale of 6 million containing semantic labels and prompts/images used to generate videos. Furthermore, GenVidBench performs cross-source setting of the training set and the test set, which is extremely challenging for AI-generated video detection.

organized into the training set and the test set, respectively, to make the dataset more challenging. Furthermore, the generators in the training and test sets are different, preventing videos generated by the same generator from sharing identical attributes. Based on this design, the task of cross-source and cross-generator detection is extremely challenging.

GenVidBench also provides rich semantic content labels for a portion of the videos. Researchers can select corresponding types of videos as required to adapt to different application scenarios. We describe the content of the video in three dimensions: object category, action, and location, which is shown in Figure 1. 1) Object categories: the main character of the given video decides the major content. 2) Actions: a reflection of the temporal attributes. 3) Locations: an indicator of scenario complexity. Through analysis, we find that the object category affects the characteristics of the generated video. Through category differentiation, researchers are able to precisely identify and concentrate on the scenes of interest.

Comprehensive experimental results are presented to establish a solid foundation for researchers working on the development and assessment of AI-generated video detectors. A number of state-of-the-art video classification models are evaluated on the GenVidBench dataset, including DeMamba, VideoSwin, UniformerV2, etc. The findings from our experiments underscore the considerable challenge inherent in cross-source and cross-generator tasks. Moreover, we conduct a detailed analysis of particularly difficult cases, employing semantic content labels to identify and filter the most demanding categories. The results from these categories are presented in detail, offering an expanded set of benchmarks to further assist researchers in the field.

Related Work

AI-Generated Content Detection Dataset

With the rapid development of generative models, the requisite expertise and effort to generate content have been reduced. This has led to a growing focus on distinguishing real items from AI-generated content and constructing corresponding detection datasets. We categorize these detection datasets into three categories: AI-generated images, deepfake videos, and AI-generated videos.

AI-Generated Image Detection Dataset. As a result of blossoming of diffusion models, AI-generated images also have become more realistic. GenImage (Zhu et al. 2024) is

a million-scale benchmark for detecting AI-generated images, which contains generated image pairs based on ImageNet (Deng et al. 2009) using various diffusion-based models and GAN-based models. Additionally, WildFake (Hong and Zhang 2024), ArtiFact (Rahman et al. 2023), and DiffusionDB (Wang et al. 2023c) also demonstrate the potential to provide a more comprehensive benchmark for fake image detection.

Deepfake Video Detection Dataset. A significant amount of research has focused on detecting deepfake videos (Ju et al. 2023; Zhang et al. 2024; Zhao et al. 2021; Ciarrarra et al. 2024; Hou et al. 2023), based on deepfake datasets such as the Deepfake Detection Challenge Dataset (DFDC) (Dolhansky et al. 2020), Celeb-DF (Li et al. 2020), FaceForensics++ (Rössler et al. 2019), and DeepFakeDetection (DFD) (google 2019). These datasets mainly use GANs, VAEs, or other swapping techniques to create fake videos. However, these datasets primarily consist of facial data and therefore lack diversity in content.

AI-Generated Video Detection Dataset. In the past few years, only a few studies have focused on detecting AI-generated videos across diverse scenarios. Generated Video Dataset (GVD) (Bai, Lin, and Cao 2024) is constructed from 11,618 video samples produced by 11 generative models. GenVideo (Chen et al. 2024a) is a large-scale AI-generated video detection dataset that collects videos from 10 different generator models for training and another 10 different generator models for testing. However, neither GVD nor GenVideo have the original prompts or images, video pairs, semantic labels, and cross-source settings, as shown in Table 1. As a result, these datasets cannot avoid the problem of similar content between the training and test sets, and they are no way to distinguish between different scenarios. Generated Video Forensics (GVF) (Ma et al. 2024) consists of video pairs from four different text-to-video models using the same prompts extracted from real videos. As illustrated in Table 1, GVF contains prompts/images, video pairs, and semantic labels. However, it is too small, with only 2.8k videos, and does not include a cross-source setting. The proposed GenVidBench is the first dataset at the scale of 6.78 million videos, containing semantic labels and prompts/images used to generate videos. Furthermore, GenVidBench adopts a cross-source and cross-generator setting between the training and test sets, making fake video detection extremely challenging.

Video Source	Train/Test	Type	Task	Resolution	FPS	V-6M	V-143k
Vript (University 2024)	Train	Real	-	-	30	417566	20131
Pika (pika 2024)	Train	Fake	T2V&I2V	1088*560	24	1670465	13501
VideoCraftV2 (Chen et al. 2024b)	Train	Fake	T2V&I2V	512*320	10	1672242	13501
Modelscope (Wang et al. 2023a)	Train	Fake	T2V	256*256	8	1672242	13501
T2V-Zero (Khachatryan et al. 2023)	Train	Fake	T2V	512*512	4	1268595	13501
OpenSora (Zheng et al. 2024)	Train	Fake	T2V	256*256	8	13800	0
HD-VG-130M (Wang et al. 2023b)	Test	Real	-	1280*720	30	13853	13853
MuseV (Xia et al. 2024)	Test	Fake	I2V	1210*576	12	13853	13853
SVD (Blattmann et al. 2023)	Test	Fake	I2V	1024*576	10	13853	13853
Mora (Yuan et al. 2024)	Test	Fake	T2V	1024*576	10	13853	13853
CogVideo (Hong et al. 2022)	Test	Fake	T2V	480*480	4	13853	13853
Sora (OpenAI 2024)	Test	Fake	T2V	1920*1080	30	51	0
Kling (KlingAI 2025)	Test	Fake	T2V&I2V	-	30	264	0
Sum	-	-	-	-	-	6784490	143400

Table 2: Statistics of real and generated videos in the GenVidBench dataset. GenVidBench contains 11 subsets of fake videos generated by 11 state-of-the-art generators and 2 subsets of real videos.

Generated Video Detections

AI-generated videos have the potential to accelerate the dissemination of misinformation, prompting significant concern. In the past, much of the research has concentrated on detecting videos with generated faces (Xu et al. 2023). However, the content of the fake face video is single, which is greatly different from the real-world scenario. Due to the lack of large-scale high-quality AI-generated video datasets, there are few work on AI-generated video detectors. Recent work on AI-generated video detection includes AIGDet (Bai, Lin, and Cao 2024), DeCoF (Ma et al. 2024). AIGDet (Bai, Lin, and Cao 2024) captures the forensic traces with a two-branch spatio-temporal convolutional neural network to improve detection accuracy. DeCoF (Ma et al. 2024) is based on the principle of video frame consistency to eliminate the impact of spatial artifacts. Beyond the specialized models dedicated to the detection of generated videos, several video classification models have demonstrated remarkable performance on this task, including VideoSwin (Liu et al. 2022) and UniFormerV2 (Li et al. 2022a). In this paper, we present the experimental results across various models on GenVidBench, which will provide a solid research foundation for developers in related fields.

Dataset Construction

Overview of GenVidBench

The GenVidBench dataset is distinguished by its large scale and exceptional diversity, providing a robust foundation for the development and evaluation of AI-generated video detection models. As illustrated in Table 2, the dataset comprises a total of 6784490 video samples. This extensive scale ensures statistical robustness and facilitates the generalization of detection models across diverse scenarios, addressing the complexities inherent in real-world applications.

The AI-generated videos in the GenVidBench dataset are categorized into two main pairs, both derived from the

same prompts or corresponding images. The first pair comprises videos from Pika, VideoCrafter2, Modelscope, and Text2Video-Zero. These videos are generated from a uniform prompt and sourced from the VidProM dataset, with each model providing approximately 1.67 million videos. The second pair includes videos from SVD, MuseV, Mora, CogVideo, and HD-VG-130M. Videos from SVD and MuseV are generated by extracting frames from the HD-VG-130M dataset. Similarly, Mora and CogVideo videos are generated using prompts from the HD-VG-130M dataset. The semantic content of the five video subsets is the same, and each model contributes 13,853 videos to the dataset. By ensuring that the content of paired videos is identical, GenVidBench effectively prevents AI video detectors from distinguishing between real and generated videos based solely on content, thereby increasing the dataset’s challenge and promoting the development of more sophisticated detection methods.

Moreover, the dataset exhibits remarkable diversity in terms of video types, tasks, resolutions, and frame rates. It encompasses both real and AI-generated content, supporting tasks such as Text-to-Video (T2V) and Image-to-Video (I2V). The resolutions range from 256×256 in lower-quality samples to high-definition 1920×1080 , while frame rates vary from 4 FPS (e.g., CogVideo) to 30 FPS (e.g., HD-VG-130M). This heterogeneity ensures comprehensive coverage of diverse content, making GenVidBench a critical benchmark for advancing the performance and generalizability of AI-generated video detection models.

Due to the substantial size of GenVidBench, the required computational resources and training time can be significant. To address this issue, we introduce GenVidBench-143k, a lightweight version of the dataset specifically designed to facilitate rapid model iteration and development. As illustrated in Table 2, this 143k subset is meticulously sampled from the original 6.78 million videos, ensuring that it remains both representative and sufficiently diverse. By focusing on this smaller yet comprehensive subset, researchers can conduct

Training Set	Test Set								Top-1
	Pika	VC2	MS	T2VZ	MuseV	SVD	Cogvideo	Mora	
Pika	96.93	95.14	54.91	55.17	65.63	54.66	69.86	53.63	68.24
VC V2	66.75	99.52	84.37	73.03	66.09	54.86	59.36	69.66	71.71
ModelScope	50.60	51.11	99.96	52.68	59.68	54.37	41.03	49.11	57.32
T2V-Zero	50.56	51.98	55.46	99.96	60.28	55.06	43.43	50.66	58.42
MuseV	70.59	66.34	54.06	55.80	99.64	62.55	73.54	53.14	66.96
SVD	50.60	61.63	62.37	65.99	94.27	99.71	80.67	81.58	74.60
Cogvideo	51.07	60.84	60.36	68.07	60.64	57.61	97.40	52.55	63.57
Mora	51.84	76.14	62.12	62.48	60.96	86.17	51.84	99.42	68.87

Table 3: Results of cross-validation on different training and testing subsets using VideoSwin-Tiny.

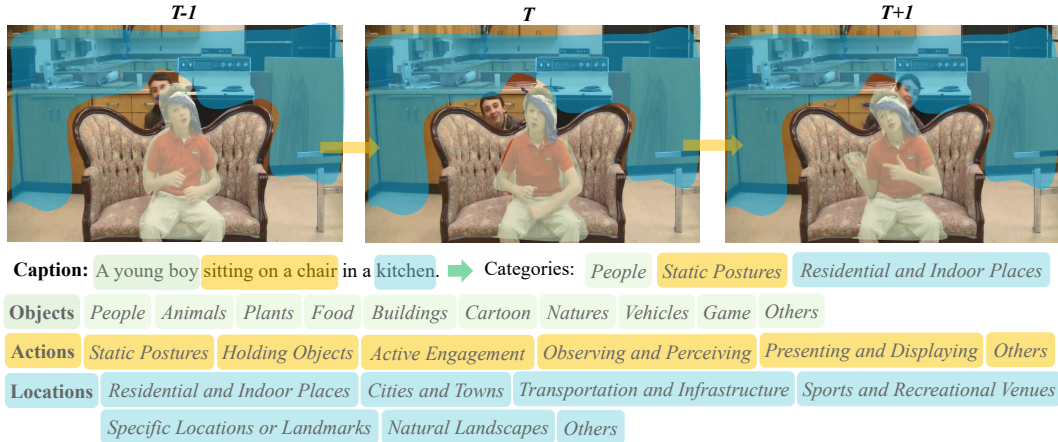


Figure 2: Semantic categorization of prompts in GenVidBench, including three dimensions: object categories, actions, and locations.

Pair1	Same	Different	Pair2	Same	Different
Pika	68.41	60.95	Musev	63.08	61.70
VC2	74.72	62.49	SVD	85.51	60.15
MS	51.46	51.05	Cogvideo	56.93	60.09
T2V-Z	52.67	52.36	Mora	65.66	63.15
Mean	61.81	56.71	Mean	67.79	61.27

Table 4: Comparison of VideoSwin-Tiny performance between training and test sets under same-source and cross-source generation. Results indicate that testing on different generation sources poses greater challenges.

in-depth analysis with significantly reduced computational demands and time overhead, while still maintaining the broad applicability.

Cross-Source and Cross-Generator Task

GenVidBench contains two sets of video pairs with identical sources. The first video pair includes Pika, ModelScope, VideoCraftV2, and T2V-Zero, while the second video pair includes MuseV, SVD, CogVideo, Mora and HD-VG-130M. Each pair uses the same images or text for generation. To an-

alyze the impact of different sources and different generators on the detection model, we carry out relevant evaluations.

First, we evaluate the performance of the detector when trained and tested on videos generated by the same generator. Our evaluation utilizes the widely recognized VideoSwin-tiny model (Liu et al. 2022). As shown in Table 3, training and testing within each subset consistently achieve accuracies above 97.40%. Among them, the T2V-Zero and ModelScope subsets perform particularly well, with accuracies exceeding 99.90%. These findings demonstrate that when the training and testing data are derived from the same generator, the task of detecting generated videos becomes relatively straightforward.

However, when testing on videos generated by a new generator, the experimental results are not satisfactory. It can be observed that a substantial performance degradation occurs when training and testing are conducted using different generators. For example, the accuracy of VideoSwin-Tiny drops to 54.66% when trained on Pika and tested on SVD. In real-world applications, the identity of the generator is typically unknown, which poses a significant challenge for reliable detection. To address this issue, our study aims to rigorously assess the generalization capability of detectors. Specifically, their ability to distinguish between real and gen-




Figure 3: Data distribution across different video sources in GenVidBench-143k. Pair1 refers to Pika, VideoCrafterV2, Mod-elscope, and Text2Video-Zero. Pair2 refers to MuseV, SVD, CogVideo, Mora, and HD-VG-130M.




Figure 4: Data distribution across different semantic categories in GenVidBench-143k.

erated videos without dependence on the generator’s identity. To this end, we introduce the cross-generator video classification task, aimed at evaluating the robustness and adaptability of detectors across diverse video sources.

Furthermore, we analyze the influence of the generation source, such as text prompts and images, on the classification of real and generated videos. The four models Pika, VideoCrafter2, Modelscope, and Text2Video-Zero use the same text prompt to generate a video, and these videos are the first video pair. The MuseV and SVD models are based on the same image for video generation, while CogVideo and Mora rely on the same text prompt, forming our second set of video pairs. As shown in Table 4, training with videos from the same video source tends to yield superior test results. For example, in the first video pair, Pika achieves an accuracy of 68.41% when tested on videos from the same source, but only 60.95% when tested on videos generated by other sources. Similarly, in the second video pair, SVD reaches 85.51% accuracy when trained and tested on the same source, yet this drops to 60.15% when evaluated on videos from different sources. Moreover, across both pairs, the average accuracy obtained with same-source testing consistently exceeds that of cross-source testing. These findings highlight the strong dependency of detectors on the generation source. Based on the above analysis, we propose a challenging dataset and introduce a cross-source and cross-generator task to enhance

detector robustness and generalization. The videos in Video Pair 1 are used as the training set, while the videos in Video Pair 2 are used as the test set. This design prevents the model from distinguishing real and fake videos merely based on content or video quality, thereby increasing the difficulty of the task and better simulating real-world conditions.

Content Analysis

Taking the 143k dataset as a representative example, we analyze the content distribution of the GenVidBench dataset. To facilitate the development of more generalized and effective detection models, we categorize prompts across multiple dimensions. As depicted in Figure 2, we initially divide prompts into three key dimensions: object categories, actions, and locations, as these three elements can form a complete sentence or story with rich semantics.

1) **Object categories.** The main subject of the given video determines the primary spatial content. 2) **Actions.** Reflect the temporal attributes of the video. For instance, categorizing the action as Static Postures implies that the main subject of the video will not exhibit rapid movement. If the action is categorized as Presenting and Displaying, the main subject may engage in walking around objects or performing other bodily movements. 3) **Locations.** This indicates the complexity of the scenario. For example, a Nature Landscape may have a clean background, whereas a Transportation scenario

Method	Type	MuseV	SVD	CogV	Mora	Sora	Kling	HD	Top-1
I3D (Carreira and Zisserman 2017)	CNN	32.72	12.04	76.44	72.3	41.18	46.22	95.04	60.21
SlowFast (Feichtenhofer et al. 2019)	CNN	87.14	29.80	93.07	55.23	23.53	58.33	96.61	70.06
TPN (Yang et al. 2020)	CNN	56.34	24.30	71.87	73.37	9.80	46.59	99.07	66.23
TIN (Shao, Qian, and Liu 2020)	CNN	68.38	41.44	57.88	79.67	35.29	60.61	96.41	67.91
TRN (Zhou et al. 2018)	CNN	43.52	27.16	58.30	97.26	29.42	58.72	97.56	66.67
TSM (Lin, Gan, and Han 2019)	CNN	95.94	73.16	36.44	91.72	33.34	71.6	96.30	73.88
TimeSformer (Bertasius, Wang, and Torresani 2021)	Trans.	69.94	29.13	67.82	98.47	64.71	65.91	72.47	71.53
UniFormerV2 (Li et al. 2022a)	Trans.	27.97	12.84	94.79	78.21	5.88	21.59	99.46	65.31
VideoSwin (Liu et al. 2022)	Trans.	90.24	27.72	91.64	88.14	19.60	50.76	99.10	80.39
MViTv2-S (Li et al. 2022b)	Trans.	77.08	44.89	99.91	76.77	61.36	31.37	86.98	80.45
DeMamba (Chen et al. 2024a)	Mamba	85.04	48.81	98.66	90.23	1.96	33.71	99.86	85.47

Table 5: Results of various state-of-the-art methods trained on cross-source and cross-generator task.

Method	Type	Top-1	F1	AUROC
I3D	CNN	60.21	67.38	84.06
SlowFast	CNN	70.06	81.13	88.65
TPN	CNN	66.23	73.39	95.67
TIN	CNN	67.91	73.74	72.96
TRN	CNN	66.67	72.49	92.13
TSM	CNN	73.88	82.49	95.31
TimeSformer	Transformer	71.53	77.31	76.71
UniFormerV2	Transformer	65.31	73.03	96.73
VideoSwin	Transformer	80.39	86.36	96.05
MViTv2	Transformer	80.45	85.66	90.29
DeMamba	Mamba	85.47	90.27	99.28

Table 6: Comparison of Top-1 accuracy, F1 score, and AUROC across different methods.

might include intricate machinery in the background.

For each dimension, taking the object category dimension as an example, we sample the entire prompt set and utilize Large Language Models (LLMs) to extract the subjects. We then aggregate these subjects into more abstract categories. The final categories are constrained to no more than ten options to construct a classification tree. For object dimension, we use people, animals, buildings, natures, plants, cartoon, food, game, vehicles, others as our final classification tree. Based on the classification tree, we can obtain the class labels for each prompt of our entire dataset.

We select object categories as our primary dimension for selecting and generating video pairs, ensuring that our benchmark is balanced and semantically rich in the spatial dimension. Figure 3 illustrates the distribution of various prompt sets across multiple dimensions. Figure 4 presents the distribution across different semantic categories. These figures demonstrate that our benchmark encompasses a wide variety of video content, spanning broad category distributions across both spatial and temporal dimensions.

Scenario-Specific Task

As mentioned above, we conduct a comprehensive content analysis of all prompts and videos to obtain category labels across multiple dimensions. This process enables researchers

Method	MuseV	SVD	CogV	Mora	HD	Top-1
I3D	8.15	8.29	60.11	59.24	93.99	45.96
SlowFast	12.25	12.68	38.34	45.93	93.63	40.57
TPN	37.86	8.79	68.25	90.04	97.34	60.46
TIN	33.78	21.47	81.59	79.44	97.88	62.83
X3D	92.39	37.27	65.72	49.60	97.51	68.50
TRN	38.92	26.64	91.34	93.98	93.97	68.97
TSM	70.37	54.70	78.46	70.37	96.76	74.13
UniFormerV2	20.05	14.81	45.21	99.21	96.89	55.23
TimeSformer	73.14	20.17	74.80	39.40	92.32	59.97
VideoSwin	49.11	68.48	88.47	81.42	98.55	77.21
MViTv2-B	76.34	98.29	47.50	96.62	97.58	83.27

Table 7: The results in GenVidBench-143k.

Dataset	SlowFast	I3D	F3Net
DeepFakes (google 2019)	97.53	-	-
Face2Face (Thies et al. 2020)	94.93	-	-
FaceSwap (Kowalski 2024)	95.01	-	-
NeuralTextures (Thies 2019)	82.55	-	-
GVF (Ma et al. 2024)	-	61.88	-
GenVideo (Chen et al. 2024a)	-	-	51.83
GenVidBench(Ours)	70.06	60.21	42.52

Table 8: Performance comparison of detectors on different datasets.

to extract datasets customized to their specific research interests. For instance, they can extract scenes of People to evaluate the logic of human body generation, scenes of Presenting to study motion deformation, and scenes with low temporal continuity to analyze frame-wise deformation. Furthermore, based on the extracted dataset, we can evaluate the impact of these attributes on the detector, such as the relationship between motion deformation and detection accuracy.

Experimental Analysis

Implementation Details

We sample 8 frames from each video as input and resize each image to 224×224 . The batch size is 8. We use the default learning rate provided by MMAAction2 (Contributors 2020) for each method. We also apply data augmentation methods

Method	People	Animals	Buildings	Natures	Plants	Cartoon	Food	Game	Vehicles	Top-1
SVD	0.757	0.728	0.777	0.736	0.739	0.537	0.769	0.732	0.795	0.739
MuseV	0.156	0.152	0.16	0.164	0.220	0.098	0.117	0.189	0.129	0.160
Mora	0.226	0.153	0.196	0.210	0.205	0.157	0.171	0.156	0.215	0.198
CogVideo	0.014	0.008	0.010	0.023	0.012	0.004	0.010	0.036	0.013	0.017
Mean	0.307	0.277	0.306	0.303	0.318	0.209	0.286	0.295	0.308	0.297

Table 9: Distribution of hard cases across categories in video subsets generated by different generators.

Method	MuseV	SVD	CogV	Mora	HD	Mean
I3D	39.18	23.27	91.98	78.38	78.42	62.15
SlowFast	81.63	29.80	75.31	19.31	73.03	55.30
TPN	43.67	20.00	85.80	86.87	94.61	64.24
TimeSformer	77.96	29.80	96.30	93.44	87.14	75.09
VideoSwin	57.96	7.35	92.59	47.88	98.76	52.86
UniFormerV2	13.88	7.76	41.98	95.75	97.93	64.76

Table 10: Results of various state-of-the-art methods trained on plants class.

such as random flipping and cropping. All other training settings follow the default configuration of MMAAction2.

Results of Cross-Source and Cross-generator Task

To establish a reliable benchmark, we evaluate the performance of multiple methods on GenVidBench-6M. As shown in Table 5, DeMamba achieves the highest accuracy with a Top-1 of 85.47%, followed by MViTv2 with 80.45%. Overall, Transformer-based models outperform CNN-based models. Among CNNs, TSM achieves the best performance with a Top-1 accuracy of 73.88%, which is 6.57% lower than MViTv2. These results demonstrate that the cross-source and cross-generator task is highly challenging, leaving substantial room for improvement in generative video detection. Furthermore, real videos are more easily distinguished from fake ones, with most classification accuracies exceeding 95.04%. Fake videos generated by Sora tend to be more challenging to classify correctly, reflecting their superior generation quality, while those produced by CogVideo are generally easier to detect due to limited temporal continuity. F1 and AUROC are also reported in Table 6.

The experimental results on GenVidBench-143k are presented in Table 7. MViTv2 still achieves a high accuracy. However, owing to the smaller training set, the overall accuracy is lower than that obtained on GenVidBench-6m.

Comparison with Advanced Datasets

To compare the challenges posed by different datasets, we conduct a comparative analysis of the performance of SlowFast, I3D, and F3Net across a range of datasets, as shown in Table 8. Following the results reported in previous work (Ma et al. 2024; Chen et al. 2024a), we additionally include new results trained on our benchmark. It can be observed that SlowFast, I3D, and F3Net achieve significantly better performance on other datasets than on GenVidBench. For instance, SlowFast reaches 82.55% accuracy on the NeuralTextures dataset, but only 70.06% on GenVidBench. Similarly, I3D

achieves 61.88% accuracy on the GVF dataset, compared with 60.21% on GenVidBench. These results demonstrate that our dataset is more challenging.

Hard Case Analysis

To further investigate the detection results, we conduct a hard-case analysis based on the experimental outcomes of VideoSwin-tiny. We define hard cases as generated videos for which the detector assigns a very low probability. The results are summarized in Table 9. While the detector demonstrates strong and consistent performance across all test videos, its effectiveness varies considerably in cross-generator scenarios. As shown in Table 9, SVD poses more challenging cases for the detector, whereas CogVideo produces almost no severely ambiguous cases.

Results on Scenario-Specific Task

Based on the analysis of hard examples, we select the Plants class and give the experimental results of various models on this class. As illustrated in Table 10, TimeSformer achieves the best accuracy, with a Top-1 of 75.09%. Additionally, TPN and UniFormerV2 also demonstrate commendable classification results, with Top-1 accuracies of 64.24% and 64.76%, respectively. Conversely, SlowFast lags behind with a Top-1 accuracy of only 55.30%. VideoSwin achieves the poorest performance, attaining a mere 52.86% in Top-1 accuracy. Experimental results also show that SVD has the lowest classification accuracy, which proves that its generation performance is the best. The above experimental results show that the classification performance of the model in a single scene is different from that in all scenes, so it is necessary to detection for different scenes. The rich semantic labels provided by GenVidBench can help analyze each scenario and provide more development ideas for developers.

Conclusion

In this paper, we introduce GenVidBench, a large-scale dataset comprising six million videos for AI-generated video detection. The dataset is characterized by its cross-source and cross-generator setup, and it provides diverse semantic content labels. We present comprehensive experimental results to establish a solid foundation for future research on the development and evaluation of AI-generated video detectors. The experiments demonstrate that GenVidBench is more challenging than existing datasets, thereby offering greater opportunities for advancing detectors.

References

- Bai, J.; Lin, M.; and Cao, G. 2024. AI-Generated Video Detection via Spatio-Temporal Anomaly Learning. *arXiv preprint arXiv:2403.16638*.
- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is Space-Time Attention All You Need for Video Understanding? In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendelevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; Jampani, V.; and Rombach, R. 2023. Stable Video Diffusion: Scaling Latent Video Diffusion Models to Large Datasets. *arXiv:2311.15127*.
- Carreira, J.; and Zisserman, A. 2017. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *CVPR*, 4724–4733.
- Chen, H.; Hong, Y.; Huang, Z.; Xu, Z.; Gu, Z.; Li, Y.; Lan, J.; Zhu, H.; Zhang, J.; Wang, W.; et al. 2024a. DeMamba: AI-Generated Video Detection on Million-Scale GenVideo Benchmark. *arXiv preprint arXiv:2405.19707*.
- Chen, H.; Zhang, Y.; Cun, X.; Xia, M.; Wang, X.; Weng, C.; and Shan, Y. 2024b. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. *arXiv preprint arXiv:2401.09047*.
- Ciamarra, A.; Caldelli, R.; Becattini, F.; Seidenari, L.; and Del Bimbo, A. 2024. Deepfake Detection by Exploiting Surface Anomalies: The SurFake Approach. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, 1024–1033.
- Contributors, M. 2020. OpenMMLab’s Next Generation Video Understanding Toolbox and Benchmark. <https://github.com/open-mmlab/mmaaction2>.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database.
- Dolhansky, B.; Bitton, J.; Pflaum, B.; Lu, J.; Howes, R.; Wang, M.; and Ferrer, C. C. 2020. The DeepFake Detection Challenge (DFDC) Dataset. *arXiv:2006.07397*.
- Feichtenhofer, C.; Fan, H.; Malik, J.; and He, K. 2019. Slow-fast networks for video recognition. In *Proceedings of the IEEE international conference on computer vision*, 6202–6211.
- google. 2019. Deepfakes dataset by Google, Jigsaw. Online. Available:<https://blog.research.google/2019/09/contributing-data-to-deepfakedetection.html>.
- Hong, W.; Ding, M.; Zheng, W.; Liu, X.; and Tang, J. 2022. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*.
- Hong, Y.; and Zhang, J. 2024. WildFake: A Large-scale Challenging Dataset for AI-Generated Images Detection. *arXiv:2402.11843*.
- Hou, Y.; Guo, Q.; Huang, Y.; Xie, X.; Ma, L.; and Zhao, J. 2023. Evading DeepFake Detectors via Adversarial Statistical Consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12271–12280.
- Ji, L.; Lin, Y.; Huang, Z.; Han, Y.; Xu, X.; Wu, J.; Wang, C.; and Liu, Z. 2024. Distinguish any fake videos: Unleashing the power of large-scale data and motion features. *arXiv preprint arXiv:2405.15343*.
- Ju, Y.; Hu, S.; Jia, S.; Chen, G. H.; and Lyu, S. 2023. Improving Fairness in Deepfake Detection. *arXiv:2306.16635*.
- Khachatryan, L.; Movsisyan, A.; Tadevosyan, V.; Henschel, R.; Wang, Z.; Navasardyan, S.; and Shi, H. 2023. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15954–15964.
- KlingAI. 2025. KlingAI. 2025-05-15.
- Kowalski, M. 2024. FaceSwap github. Online. Available:<https://github.com/MarekKowalski/FaceSwap>.
- Li, K.; Wang, Y.; He, Y.; Li, Y.; Wang, Y.; Wang, L.; and Qiao, Y. 2022a. Uniformerv2: Spatiotemporal learning by arming image vits with video uniformer. *arXiv preprint arXiv:2211.09552*.
- Li, Y.; Wu, C.-Y.; Fan, H.; Mangalam, K.; Xiong, B.; Malik, J.; and Feichtenhofer, C. 2022b. MVITv2: Improved multi-scale vision transformers for classification and detection. In *CVPR*.
- Li, Y.; Yang, X.; Sun, P.; Qi, H.; and Lyu, S. 2020. Celeb-DF: A Large-scale Challenging Dataset for DeepFake Forensics. *arXiv:1909.12962*.
- Lin, J.; Gan, C.; and Han, S. 2019. TSM: Temporal Shift Module for Efficient Video Understanding. *arXiv:1811.08383*.
- Liu, Z.; Ning, J.; Cao, Y.; Wei, Y.; Zhang, Z.; Lin, S.; and Hu, H. 2022. Video swin transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3202–3211.
- Ma, L.; Zhang, J.; Deng, H.; Zhang, N.; Liao, Y.; and Yu, H. 2024. DeCoF: Generated Video Detection via Frame Consistency. *arXiv:2402.02085*.
- OpenAI. 2024. Sora. Online. Available:<https://openai.com/index/sora/>.
- pika. 2024. Pika. Online. Available:<https://pika.art/home>.
- Rahman, M. A.; Paul, B.; Sarker, N. H.; Hakim, Z. I. A.; and Fattah, S. A. 2023. ArtiFact: A Large-Scale Dataset with Artificial and Factual Images for Generalizable and Robust Synthetic Image Detection. *arXiv:2302.11970*.
- Rössler, A.; Cozzolino, D.; Verdoliva, L.; Riess, C.; Thies, J.; and Nießner, M. 2019. FaceForensics++: Learning to Detect Manipulated Facial Images. *arXiv:1901.08971*.
- Shao, H.; Qian, S.; and Liu, Y. 2020. Temporal Interlacing Network. *AAAI*.
- Thies, J.; Zollhöfer, M.; Stamminger, M.; Theobalt, C.; and Nießner, M. 2020. Face2Face: Real-time Face Capture and Reenactment of RGB Videos. *arXiv:2007.14808*.
- Thies, J. e. a. 2019. Deferred neural rendering: Image synthesis using neural textures. *Acm Transactions on Graphics (TOG)*, 38(4): 1–12.
- University, S. J. T. 2024. Vript: A Video Is Worth Thousands of Words. Online. Available:<https://github.com/mutonix/Vript>.

Wang, J.; Yuan, H.; Chen, D.; Zhang, Y.; Wang, X.; and Zhang, S. 2023a. Modelscope text-to-video technical report. *arXiv preprint arXiv:2308.06571*.

Wang, W.; Yang, H.; Tuo, Z.; He, H.; Zhu, J.; Fu, J.; and Liu, J. 2023b. VideoFactory: Swap Attention in Spatiotemporal Diffusions for Text-to-Video Generation. *arXiv preprint arXiv:2305.10874*.

Wang, Z. J.; Montoya, E.; Munechika, D.; Yang, H.; Hoover, B.; and Chau, D. H. 2023c. DiffusionDB: A Large-scale Prompt Gallery Dataset for Text-to-Image Generative Models. arXiv:2210.14896.

Xia, Z.; Chen, Z.; Wu, B.; Li, C.; Hung, K.-W.; Zhan, C.; He, Y.; and Zhou, W. 2024. MuseV: Infinite-length and High Fidelity Virtual Human Video Generation with Visual Conditioned Parallel Denoising. *arxiv*.

Xu, Y.; Liang, J.; Jia, G.; Yang, Z.; Zhang, Y.; and He, R. 2023. TALL: Thumbnail layout for deepfake video detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22658–22668.

Yang, C.; Xu, Y.; Shi, J.; Dai, B.; and Zhou, B. 2020. Temporal Pyramid Network for Action Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Yuan, Z.; Chen, R.; Li, Z.; Jia, H.; He, L.; Wang, C.; and Sun, L. 2024. Mora: Enabling Generalist Video Generation via A Multi-Agent Framework. arXiv:2403.13248.

Zhang, X.; Yi, J.; Wang, C.; Zhang, C. Y.; Zeng, S.; and Tao, J. 2024. What to Remember: Self-Adaptive Continual Learning for Audio Deepfake Detection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(17): 19569–19577.

Zhao, H.; Zhou, W.; Chen, D.; Wei, T.; Zhang, W.; and Yu, N. 2021. Multi-attentional Deepfake Detection. arXiv:2103.02406.

Zheng, Z.; Peng, X.; Yang, T.; Shen, C.; Li, S.; Liu, H.; Zhou, Y.; Li, T.; and You, Y. 2024. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*.

Zhou, B.; Andonian, A.; Oliva, A.; and Torralba, A. 2018. Temporal Relational Reasoning in Videos. *European Conference on Computer Vision*.

Zhu, M.; Chen, H.; Yan, Q.; Huang, X.; Lin, G.; Li, W.; Tu, Z.; Hu, H.; Hu, J.; and Wang, Y. 2024. Genimage: A million-scale benchmark for detecting ai-generated image. *Advances in Neural Information Processing Systems*, 36.