

# Data Collection-free Masked Video Modeling

Yuchi Ishikawa<sup>1,2</sup>, Masayoshi Kondo<sup>1</sup>, and Yoshimitsu Aoki<sup>2</sup>

<sup>1</sup> LY Corporation

{yuchi.ishikawa, masayoshi.kondo}@lycorp.co.jp

<sup>2</sup> Keio University

aoki@elec.keio.ac.jp

**Abstract.** Pre-training video transformers generally requires a large amount of data, presenting significant challenges in terms of data collection costs and concerns related to privacy, licensing, and inherent biases. Synthesizing data is one of the promising ways to solve these issues, yet pre-training solely on synthetic data has its own challenges. In this paper, we introduce an effective self-supervised learning framework for videos that leverages readily available and less costly static images. Specifically, we define the Pseudo Motion Generator (PMG) module that recursively applies image transformations to generate pseudo-motion videos from images. These pseudo-motion videos are then leveraged in masked video modeling. Our approach is applicable to synthetic images as well, thus entirely freeing video pre-training from data collection costs and other concerns in real data. Through experiments in action recognition tasks, we demonstrate that this framework allows effective learning of spatio-temporal features through pseudo-motion videos, significantly improving over existing methods which also use static images and partially outperforming those using both real and synthetic videos. These results uncover fragments of what video transformers learn through masked video modeling.

**Keywords:** Self-supervised Learning · Masked Video Modeling · Action Recognition · Pseudo-motion Videos




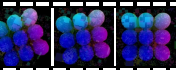
## 1 Introduction

Pre-training video transformers [3, 9] generally requires a large amount of labeled data. Although self-supervised learning enables pre-training of video transformers without labels [17, 56], it still demands substantial volumes of video data. This highlights various issues related to real video data including the following:

**High cost of data collection.** Video data, compared to audio, text, and images, is massive in size. Therefore, downloading, storing, and pre-processing videos is extremely costly. Furthermore, the following issues related to licenses, privacy, and bias arise during data collection.

**Copyright and license infringement.** Video data may have been collected without permission, potentially infringing on licenses and copyrights. For example, some datasets are gathered from video-sharing sites like YouTube,

**Table 1: Comparison of each data source and their issues when conducting pre-training.** While real videos enhance model performance, they have concerns related to collection cost, privacy and licenses. Synthetic videos [28, 33, 74] and pseudo motions by MoSI [25] partially resolve these issues, but they rely on the CNN architecture and its inherent inductive bias, thus failing to accurately train ViT. Note that in VPN, additional real data is required for optimal performance, therefore the asterisked issues (\*) are not resolved. Our proposed framework is free of these issues by generating pseudo-motion videos from synthetic images.

	Real Video [56]	VPN [28]	MoSI [25]	Ours
example				
Acc. @UCF101	96.1	89.9	82.8	89.4
collection cost		✓*	✓	✓
privacy/license		✓*	✓	✓
training ViT	✓			✓

in which videos are licensed by default with a Standard YouTube license<sup>3</sup>, which prohibits the download of content.

**Privacy issues.** Video data often contains Personally Identifiable Information (PII) including faces, which raises significant privacy concerns.

**Bias and ethical issues.** Large-scale datasets may unintentionally include biases leading to ethical issues related to nationality, gender, age, and more [10, 63, 68], which can impact the fairness and inclusiveness of model outcomes. Some works have also reported that video recognition models might have context and object biases, failing to recognize actions accurately [13, 37, 38].

**Data access issues.** Possibly due to the above issues, some datasets like IG-Curated/Uncurated dataset [18, 21] and CREATE [39, 73] are only made available to certain research groups. This limitation restricts other researchers from replicating or further developing these works, thereby impeding scientific progress.

In image recognition, to address these concerns and eliminate the costs associated with data collection, some researchers have proposed pre-training methods using synthetic images as an alternative to those using real images. While some works have synthesized images from mathematical formulas [30, 44, 54], others have utilized structured noise [7] or OpenGL fragment shaders [6]. These methods have achieved comparable results to pre-training on real image datasets like ImageNet [15] and JFT-300M [52], emphasizing the importance of data diversity.

However, pre-training using synthetic videos still presents significant challenges. Few works address this, including the Video Perlin Noise (VPN) dataset [28] generated from Perlin Noise [46, 47], and SynAPT [33, 74]. However, they still require real video datasets such as Kinetics400 [31]. This diverges from our goal of reducing data collection costs and minimizing issues related to real data.

An alternative approach involves generating pseudo-motions from static images. Huang et al. [25] proposed a self-supervised learning framework named Un-

<sup>3</sup> <https://www.youtube.com/t/terms>

masked MoSI, designed to make models learn spatio-temporal features through the classification of pseudo motions. This can be promising because it only requires static images and can be combined with datasets with protected privacy and liberal licenses, like PASS [4]. However, this method is specialized on CNN architectures and cannot generalize to transformer-based architectures, which are the current state-of-the-art models.

In this paper, to mitigate video collection costs and address concerns regarding privacy, bias, and licenses, we propose a self-supervised learning framework for video transformers using synthetic images (Table 1). Our framework includes a Pseudo Motion Generator (PMG) module that recursively applies image transformations to static images, generating videos with diverse pseudo-motion. These videos are then used for masked video modeling. Through experiments, by using videos generated from the PMG module, we examine that video transformers can learn transferable and robust video features which are not limited to a single domain. To the best of our knowledge, we are the first to pre-train video transformers exclusively using synthetic images. Our contributions are threefold;

1. We introduce a self-supervised learning framework for videos that uses single images to reduce data collection costs compared to videos. Our framework includes a Pseudo Motion Generator (PMG) module, which generates a wide variety of pseudo-motion videos. These pseudo-motion videos are utilized for self-supervised masked video modeling. Notably, PMG can also be used for video augmentation when pre-training with real videos.
2. We demonstrate that synthetic images can be used for our framework to still effectively pre-train video transformers, completely eliminating the need for real videos or images. This mitigates privacy, bias, and licensing concerns.
3. Through experiments in action recognition tasks, we demonstrate that our proposed framework significantly improves over existing works using static images, and also partially surpasses existing pre-training methods using both real and synthetic videos. These experimental findings reveal pieces of what video transformers learn through masked video modeling.

## 2 Related Work

**Self-supervised Learning for Videos.** Videos require significantly more effort than images and text for annotation. Therefore, more interest is invested in self-supervised learning methods which do not require labeled data. While earlier works leverage pretext tasks [2, 8, 19, 27, 32, 59, 66, 69], recent advancements have introduced contrastive learning [18, 45] and masked video modeling [5, 35, 53, 60–62, 64, 67], which offers more robust representation learning without explicit labeling. Notably, VideoMAE [17, 56] has emerged as a leading method due to its simplicity and efficacy, learning video representations by simply reconstructing masked regions. Some works, however, point out that VideoMAE predominantly learns low-level features such as shapes. This tendency may limit its ability to capture high-level semantic features [35, 49]. Nonetheless, the emphasis on low-level features suggests that VideoMAE does not specialize in domain-specific

features, leading to its high transferability across various domains. We aim to capitalize on this characteristic to train video transformers with static images.

**Large-scale Datasets in Computer Vision.** Though self-supervised learning eliminates the need for annotation, it still demands large volumes of data. The growth of computer vision has relied on massive datasets like ImageNet [15] and LAION-5B [48]. However, these resources are fraught with privacy, bias, and licensing issues. Furthermore, access to datasets like JFT-300M [52], Instagram-3.5B [40], IG-Curated/Uncurated [18, 21], and CREATE [39, 73] is restricted to certain research groups. These issues underscore the urgent need for accessible data sources which are free from bias and privacy violation.

Video data exacerbates these challenges with its higher collection costs, privacy risks and biases [13, 37, 38]. Some popular datasets such as Kinetics400 [31], HowTo100M [41], YouTube-8M [1], and ActivityNet [11], are collected from YouTube and may encounter copyright and license restrictions. On the other hand, our self-supervised framework requires only synthetic images which are free from these challenges.

**Learning from Synthetic Data.** In response to these challenges, there is a growing interest in synthetic data, which bypasses many of the issues existent when using real-world data. Some research has focused on synthesizing realistic data [14, 20, 26, 42, 55, 58, 65], while others have proposed systematic synthesis of data from noise [7] or mathematical formulas [6, 30, 43, 44, 54]. These works have proven that not only realism but also diversity in synthetic data is crucial for effectively training models.

Few attempts are made to train action recognition models using synthetic data. For example, the GATA dataset [23], collected from a video game, is proposed for human motion representation learning. However, this dataset is not allowed for commercial use, and the rights of game companies have not been considered. Another example is the Video Perlin Noise (VPN) dataset [28], which is generated from Perlin Noise [46, 47]. This dataset is proposed to initialize model weights before pre-training. Zhong et al. [74] propose a pre-training method with both No-Human Kinetics (NH-Kinetics) and SynAPT [33]. While these approaches contribute to model performance, they still require pre-training on real videos. Additionally, ElderSim [26], PHAV [14], and SURREAL [58], which are included in SynAPT, are not allowed for commercial use. As an alternative, Huang et al. [25] have proposed MoSI, which pre-trains models with pseudo-motion videos generated from static images. In terms of collection cost, requiring only static images for pre-training is favorable. However, because MoSI’s synthesized videos lack diversity, they fail to pre-train video transformers (See Tab. 10).

Overall, existing works have shown the capability to pre-train video recognition models using synthetic or pseudo-motion videos. However, they either specialize on CNN architectures or still require the use of real video data. In contrast, our method generates a diversity of pseudo-motion videos from synthetic images, which can effectively pre-train video transformers. Moreover, our approach is completely agnostic of the issues associated with video data collection, privacy, and bias.

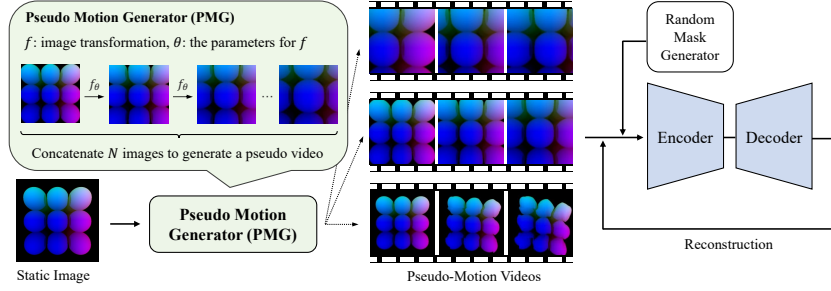


Fig. 1: Overview of our proposed framework.

### 3 Proposed Method

#### 3.1 Overview of Our Self-supervised Framework

To reduce the collection cost of video data, we propose a self-supervised framework using pseudo-motion videos generated from static images. Figure 1 shows the overview of our framework. We first generate pseudo-motion videos from static images by Pseudo Motion Generator (PMG). Then, we utilize these videos to train VideoMAE [17, 56]. VideoMAE is a powerful self-supervised learning framework and can learn spatio-temporal features effectively by reconstructing masked video regions from their complementaries. Some works point out that VideoMAE has a tendency to learn low-level features such as edges, thus failing to achieve high-level alignment [35, 49]. Conversely, VideoMAE does not obtain domain-specific features, leading to high transferability. We focus on and leverage this characteristic to train video transformers with pseudo-motion videos.

#### 3.2 Pseudo Motion Generator (PMG)

As mentioned, VideoMAE learns low-level features such as edges in a video. Especially, we hypothesize that it focuses on the correspondence of patches between frames. Therefore, we assume that to train VideoMAE effectively, patches between frames in videos should be trackable. To generate such a pseudo-motion video  $V \in \mathbb{R}^{C \times T \times H \times W}$ <sup>4</sup>, we propose a simple module, namely Pseudo Motion Generator (PMG). The algorithm of PMG is as follows: First, PMG randomly selects an image transformation  $f$  from a predefined set  $\phi$  and determines its intensity parameter  $\theta$ . Then, PMG takes as input a static image  $I_1 \in \mathbb{R}^{C \times H \times W}$  and recursively applies image transformation to  $I_1$ .

$$I_{i+1} = f_{\theta}(I_i) \quad \text{for } i = 1, \dots, T-1 \quad (1)$$

Finally, by concatenating the images from  $I_1$  to  $I_T$  in the temporal dimension, a pseudo-motion video  $V$  is generated.

$$V = [I_1; I_2; \dots; I_T] \quad (2)$$

<sup>4</sup>  $T$  is the number of frames in a video.  $H$  and  $W$  are the width and the height of each video frame.  $C$  is the number of channels.



**Fig. 2: Examples of pseudo-motion videos.** Images are sampled from PASS [4]. For more examples of pseudo-motion videos, see the supplementary material.

For clarity, we also provide pseudo-code for PMG in the supplementary material.

As candidates for image transformation, we consider the following 8 image transformations. (See Sec. 4.1 for the effect of each transformation)

1. **Identity:** Return an input image as is. We regard this as a baseline.
2. **Sliding Window:** Cut a window from a static image and move it randomly. Note that this is similar to Unmasked MoSI [25], but our method does not limit the window’s movement to only four directions.
3. **Zoom-in/out:** Cut a window from an input image and enlarge or reduce the size of the window.
4. **Fade-in/out:** An input image gradually becomes visible or invisible.
5. **Affine Transformation**
6. **Perspective Transformation**
7. **Color Jitter:** Randomly change the brightness, contrast, saturation, and hue of an input image.
8. **CutMix:** Generate an image using CutMix [70] from two images and move a small area of the image in the manner of Sliding Window.

From these candidates, through experimentation, we identify the optimal set of image transformations  $\phi$  (See Sec. 4.1). Furthermore, to prevent overfitting to specific types of pseudo-motion videos, we apply mixup [72] to each frame of the generated pseudo-motion videos. This approach significantly enhances the diversity in motion and appearance of the pseudo-motion videos, facilitating more efficient learning by VideoMAE. In the supplementary material, we describe the parameters for each image augmentation.

Figure 2 presents examples of pseudo-motion videos generated by PMG. Although the motions in these videos differ from real videos, they exhibit a wide range of motion and appearance patterns. Moreover, the clear correspondence of patches between frames makes these pseudo-motion videos particularly well-suited for VideoMAE, because it focuses on capturing low-level features rather than high-level semantic features. Notably, when pre-training VideoMAE using real videos, we can use pseudo-motion videos generated from a frame within the videos as a powerful form of data augmentation (we call this PMG Aug). We demonstrate the effect of PMG Aug through experiments (Refer to Sec. 4.4).

### 3.3 Combination of Our Framework with Synthetic Images

Our framework enables the pre-training of video transformers using single images, which are more accessible than real videos. Additionally, our framework is applicable to synthetic images, further reducing data collection costs and minimizing privacy and other concerns associated with the use of real-world data. We use the following synthetic image datasets for this purpose: (i) FractalDB, generated based on fractal geometry [30], (ii) Visual Atom, created using sine waves [54], (iii) Shaders1k, produced through OpenGL fragment shaders [6]. These datasets encompass a large volume and wide variety of images, and have demonstrated to be as effective as real image datasets in the image recognition task. By combining these synthetic images with our PMG module, we can generate a wide variety of pseudo-motion videos, enabling video transformers to learn effective spatio-temporal representations as they would using real videos.

## 4 Experiments

**Datasets.** To evaluate the effectiveness of our framework, we pre-train on pseudo-motion videos before fine-tuning and evaluating on various action recognition datasets. Following the SynAPT benchmark, we use six action recognition datasets for fine-tuning and evaluation; UCF101 [50], HMDB51 [34], MiniSSV2 [12] (a subset of Something-Something V2 [22]), Diving48 [38], IkeaFA [57], and UAV-Human (UAV-H) [36]. This benchmark is used to assess the transferability of our framework. Additionally, we use Kinetics400 (K400) [31]. As an evaluation metric, we report the top-1 accuracy.

For pre-training, we adopt randomly sampled images from the following large-scale image datasets: ImageNet-1k (IN-1k) [15], PASS [4], FractalDB [30], Shaders1k [6], and Visual Atom [54]. If the datasets have category annotations, we sampled images so that the number of images of each category is the same. Additionally, for comparison with [25], we randomly sample one frame of a video and use it as an input image for generating pseudo-motion videos.

**Implementation Details.** We conducted our experiments using 8 A100 GPUs. Our training settings were mostly aligned with VideoMAE [56], with a mask ratio of 0.75 and the number of epochs set to 2,000 unless otherwise noted (See the supplementary material for details). We used videos with 16 square frames (224 pixels in width). For the model architecture, we adopted a vanilla ViT [16] as the backbone, specifically the ViT-Base variant.

### 4.1 Ablation Studies

**The effect of image augmentations.** First, we investigated the contribution of each image transformation on VideoMAE pre-training. We used HMDB51 and UCF101 for generating pseudo-motion videos for pre-training, then used videos from the respective datasets to fine-tune the model. Tab. 2 reports the results when applying only a single variation of image augmentation. Videos generated by the Identity transformation serve as a baseline because they do not contain any motion. Compared to this baseline, videos generated with Sliding Window,

**Table 2: Comparison of different image augmentations.**

Method	UCF101	HMDB51
Baseline (Identity)	72.7	35.6
Sliding Window	75.1	40.5
Zoom-in/out	81.2	44.5
Fade-in/out	76.3	34.1
Affine	80.5	43.2
Perspective	<b>82.7</b>	<b>45.9</b>
Color Jitter	76.2	38.7
CutMix	76.8	45.1

**Table 3: Combination of image augmentations for PMG.**

Zoom-in/out	Affine	Perspective	CutMix	HMDB51
✓	✓			<b>51.8</b>
✓		✓		45.2
✓			✓	41.4
	✓	✓		50.5
	✓		✓	47.1
		✓	✓	44.2
✓	✓	✓		49.0
✓	✓		✓	49.4
✓		✓	✓	47.2
	✓	✓	✓	42.0
✓	✓	✓	✓	47.9

Zoom-in/out, Affine Transformation, Perspective Transformation, and CutMix improve the model’s accuracy over the baseline. Pseudo-motion videos generated with these transformations have corresponding patches between frames, meaning that patches in one frame might slightly move but would still exist in the subsequent frame. Therefore, this supports our hypothesis that this characteristic aids the VideoMAE when learning spatio-temporal features.

While videos generated with Fade-in/out and Color Jitter marginally improved performance on UCF101, they did not do as well on HMDB51, which is a motion-sensitive dataset [38]. This suggests that videos made with these transformations are beneficial for capturing spatial features but do not aid in capturing motion features. Next, we experimentally determine the optimal set of image transformations  $\phi$  from Sliding Window, Zoom-in/out, Affine Transformation, Perspective Transformation, and CutMix.

**The combination of image augmentations.** Tab. 3 compares the performance on HMDB51 when models are pre-trained with various combinations of image transformations. It is observed that combining multiple image transformations improves the model’s performance. This indicates that the model can effectively learn as long as there is sufficient diversity, even if the motion patterns in pseudo-motion videos differ from those in real videos. However, combining more image transformations did not necessarily yield better results. In particular, in most cases where we applied CutMix, the accuracy decreased. We hypothesize that this is due to the non-continuous nature of CutMix videos. From this point on, we will use Zoom-in/out and Affine Transformation as the set of image transformations  $\phi$ . Further discussion on the failure cases of pre-training with these pseudo-motion videos is provided in the supplementary material.

**The efficacy of video-level augmentations.** To further enhance the diversity of videos, we applied video-level augmentation to the generated pseudo-motion videos. We examined two methods: Mixup [72] and VideoMix [71]. Tab. 4 demonstrates that video-level augmentation, especially Mixup, significantly contributes to performance improvement. This is because both video augmentations diversify pseudo-motion videos, resulting in better performance. Pre-training with VideoMix results in lower accuracy compared to Mixup because the videos



**Table 4: Effects of video-level augmentation.**

Video Augmentation	Mixup	Dataset	
		HMDB51	UCF101
		51.8	83.8
✓		<b>55.9</b>	<b>87.3</b>
	✓	53.0	85.2

**Table 5: Transferability from other video datasets.**

Pre-training	Fine-tuning	Top1
HMDB51	HMDB51	55.9
UCF101	HMDB51	56.7
UCF101	UCF101	87.3
HMDB51	UCF101	85.5

generated by VideoMix have non-continuous regions like CutMix, as discussed in Sec. 4.1, From here, we will utilize Mixup in our experiments.

## 4.2 Transferability of Our Framework

**Transferability from other video datasets.** To verify the transferability of our framework, we conducted experiments by pre-training models with pseudo-motion videos generated from frames in HMDB51 and then fine-tuning on UCF101 (hereafter, we refer to this as HMDB51  $\rightarrow$  UCF101), and then vice versa (UCF101  $\rightarrow$  HMDB51). Tab. 5 shows the results. Comparing the accuracy when pre-training on different datasets, the difference is marginal. This suggests that our framework learns robust features that are not domain-specific. Furthermore, this appeals that our framework can effectively pre-train models even when using image datasets instead, such as ImageNet and PASS.

**Transferability from real image datasets.** In our previous experiments, we used samples with similar visual cues between pre-training and fine-tuning, namely the semantic information including objects and people. To further assess the transferability of our framework, we conducted pre-training on the ImageNet-1k and PASS, which are in different domains compared to the fine-tuning datasets (UCF101 and HMDB51). As detailed in Tab. 6, pre-training using ImageNet and PASS achieved comparable performance to when pre-training with the same datasets that are used when fine-tuning. Note that PASS does not include any human images. Therefore, the semantic information within pre-training datasets are not a must for effective pre-training of VideoMAE. Moreover, increasing the number of images scaled the performance. These experimental results suggest that for VideoMAE, the diversity of the data is more crucial than domain-specific information like human motion or visual cues.

**Transferability from synthetic images.** We then pre-trained on synthetic image datasets using our framework to verify that spatio-temporal features can be effectively learnt from synthetic images, which present completely different visual cues compared to our target action recognition datasets. For synthetic image datasets, we used FractalDB, Shaders1k, and Visual Atom. Herein, we used 10k/100k images sampled from each dataset. Tab. 7 shows the performance on UCF101 and HMDB51 when pre-training on diverse synthetic datasets, including FractalDB, Shaders1k, and VisualAtom. Note that pre-training with Shaders1k achieved comparable results to pre-training with real images, where pre-training with FractalDB and Visual Atom lead to subpar performance. This denotes

**Table 6: Pre-training with ImageNet and PASS.** The term 'FT data' indicates that the datasets used for pre-training are identical to those used in fine-tuning.

Pre-training Dataset	Pre-training #data	Downstream task	
		UCF101	HMDB51
FT data	-	87.3	55.9
ImageNet	10,000	87.4	58.0
ImageNet	100,000	89.2	59.2
PASS	10,000	87.6	58.3
PASS	100,000	<b>89.3</b>	<b>60.0</b>

**Table 7: Pre-training on synthetic image datasets.**

Pre-training Dataset	Setting #Data	Downstream task	
		UCF101	HMDB51
FT data	-	87.3	55.9
FractalDB	10,000	77.6	42.8
	100,000	78.1	41.1
Shaders1k	10,000	88.4	57.6
	100,000	<b>89.6</b>	<b>59.7</b>
Visual Atom	10,000	83.5	48.9
	100,000	82.6	48.2

that the model struggles to correlate patches between frames of pseudo-motion videos generated from FractalDB and Visual Atom, thus failing to capture robust low-level features. On the other hand, images in Shaders1k have distinctive patches that can be correlated before and after transformations, which supports the model when capturing low-level features. This indicates that our framework can successfully replace the need for real data when pre-training the model, as long as synthetic videos have patches that can be tracked between frames. Thus, when using our framework, challenges related to real datasets such as privacy and license infringement are nonexistent.

### 4.3 Effect of the Number of Epochs, Data, and Categories in Image Datasets for Pre-training

Fig. 4a shows the relationship between the number of epochs of pre-training and accuracy on HMDB51. For generating videos for pre-training, we used 10k images from Shaders1k and a frame from each of the 3k videos in HMDB51. In both datasets, the model performance improved over epochs and the difference of accuracy gradually decreased. Because our PMG allows for the generation of diverse videos, even if we have a small amount of data for pre-training, it is possible to improve performance by increasing the number of iterations.

Fig. 3b presents the accuracy transition when the number of pre-training samples is varied among  $\{1k, 5k, 10k, 50k, 100k\}$ . Our framework shows improvement as the number of data increased. Because we use only a small subset from PASS and Shaders1k, there is potential for more substantial performance improvement when generating from all images.

Based on the results of the previous experiment, we hypothesized that performance can be further enhanced by increasing the diversity of samples in pre-training image datasets. Fig. 4b shows the relationship between the number of categories in the pre-training datasets we use and the classification performance on HMDB51. We set the number of training samples to 10k images, using the IN-1k and Shaders1k datasets. For IN-1k, the accuracy seems to saturate after raising the diversity to more than 50 classes. For Shaders1k, the accuracy was almost the same even when the number of categories increased. This suggests our framework scales with having more data samples, but does not require




Fig. 3: Effect of the number of epochs, data, categories.

Table 8: Effectiveness of our PMG as a video augmentation method on HMDB51 and UCF101. \* results from [56].

Pre-training data		Downstream tasks	
real video	pseudo motion	HMDB51	UCF101
✓		62.6*	91.3*
	✓	55.9	87.3
✓	✓	<b>64.6</b>	<b>92.2</b>

Table 9: Comparison of each combination with real videos and pseudo-motion videos. \* Sources of PMG Aug.

Pre-training data				HMDB51
Videos	frames*	PASS*	Shaders1k*	
✓				62.6
✓	✓			64.6
✓	✓	✓		<b>68.0</b>
✓	✓		✓	67.0
✓	✓	✓	✓	67.9

semantic diversity within the samples. These results support the fact that VideoMAE learns low-level features like the correspondence of patches between frames, rather than semantic information like categories of objects displayed.

#### 4.4 PMG as Video Augmentation on Pre-training

**Pre-training with both real videos and pseudo-motion videos.** As we have verified so far, our framework enables efficient pre-training with static images. This suggests that our proposed PMG can be also used as a data augmentation method during pre-training. Tab. 8 compares the performance when pre-training solely with real videos and when pre-training with both real and pseudo-motion videos. Notably, PMG Aug boosted model accuracy by up to 2%. This suggests that synthetic motion, despite its differences to real video motion, unintuitively contributes to the model’s performance by increasing diversity.

**Can we combine image datasets with video datasets to train our framework?** Next, we use image datasets as well as sources of PMG Aug during pre-training. For the image datasets (PASS and Shaders1k), we randomly sampled 10k images as input. Tab. 9 compares the performance of the models pre-trained on HMDB51, PASS, and Shaders1k. The results show that using both image and video datasets improved the model’s performance. Particularly, the combination of HMDB51 and PASS enhanced the accuracy on HMDB51 by 5.4% compared to pre-training with only real videos. This indicates that using PMG Aug resolves the problem of insufficient data quantity during VideoMAE pre-training.

#### 4.5 Comparison to Existing Methods

**Comparison to methods using HMDB51, UCF101 and Diving48.** The upper part of Tab. 10 presents the performance of existing works which pre-train

**Table 10: Comparison with existing methods on HMDB51, UCF101, and Diving48.** RV = Real Videos, SV = Synthetic Videos, RI = Real Images, SI = Synthetic Images, SP = Supervised Pre-training, FT data = Fine-tuning data. <sup>†</sup> Results in our replicated experiments. <sup>‡</sup> Reported in [33]. \* Herein, we refer to a combination of ElderSim [26], SURREACT [58], and PHAV [51] as SynAPT, as proposed in [33]. <sup>§</sup> we report only the number of videos in SynAPT.

Method	Pre-training Setting			Downstream Tasks		
	Dataset	Data Source	#Data	UCF101	HMDB51	Diving48
from scratch (ViT-B)	-	-	-	51.4	18.0	17.9 <sup>†</sup>
MoCo v3 (ViT-B) [12]	FT data	RV	-	81.7	39.2	-
VideoMAE (ViT-B) [56]	FT data	RV	-	91.3	62.6	79.3 <sup>†</sup>
VideoMAE (ViT-B) [56]	Kinetics400	RV	260k	96.1	73.3	-
VideoMAE (ViT-B) <sup>†</sup>	VPN [28]	RV	10k	64.9	30.3	17.5
3D-ResNet50 [24]	VPN [28]	SV	28k	49.9	23.0	-
3D-ResNet50 [24]	VPN→Kinetics400	RV + SV	280k	89.9	61.8	-
TimeSformer [9] <sup>‡</sup>	IN-21k→SynAPT*	RI + SV	150k <sup>§</sup>	89.0	54.4	44.9
PPMA [74]	NH-Kinetics+SynAPT*	RV + SV	300k	92.5	71.2	64.0
MoSI (R-2D3D) [25]	FT data	RI	-	71.8	47.0	-
MoSI (R(2+1)D) [25]	FT data	RI	-	82.8	51.8	-
MoSI (ViT-B) <sup>†</sup>	FT data	RI	-	48.0	27.3	14.2
SP (ViT-B) <sup>†</sup>	IN-21k	RI	14M	71.9	34.0	34.2
SP (ViT-B) <sup>†</sup>	ExFractalDB-21k [29]	SI	21M	61.5	20.8	28.0
SP (ViT-B) <sup>†</sup>	VisualAtom-21k	SI	21M	58.9	20.3	21.4
Ours (ViT-B)	frames from FT data	RI	-	87.3	55.9	68.3
	PASS	RI	100k	89.3	<b>60.0</b>	69.2
	Shaders1k	SI	100k	<b>89.4</b>	59.7	<b>72.3</b>

using the HMDB51, UCF101, and Diving48 datasets. Existing methods like 3D-ResNet with VPN [28], TimeSformer with SynAPT, and PPMA have improved model performance compared to training the model from scratch. However, they still require real data, causing issues as mentioned. In contrast, our framework, despite using fewer samples which are also synthetic, achieves comparable performance on UCF101 and better performance on Diving48.

We also compared with pre-training methods which only use static images (the lower part of Tab. 10). MoSI works on CNN-based architectures, but it fails to pre-train a ViT model because of the lack of diversity in generated videos. Supervised pre-training (SP) on IN-21k, ExFractalDB-21k [29] and VisualAtom-21k slightly improves the performance in comparison with 'from scratch'. However, our framework significantly surpasses that performance in both settings, when using real images and when using synthetic images.

Note that VideoMAE pre-trained with VPN has low accuracy on downstream classification tasks, which suggests that VPN does not work well with VideoMAE when learning spatio-temporal features. We consider this is because VPN videos have temporal continuity, but do not possess clear correspondence of patches between frames (e.g. edges are ambiguous, and regions suddenly disappear or appear). We believe this characteristic is key for effective VideoMAE pre-training. In Sec. 4.7, we further experiment to support this hypothesis.

**Comparison on SynAPT benchmark.** Following the SynAPT benchmark [33], we evaluate using the following six datasets: UCF101, HMDB51, MiniSSV2, Div-

**Table 11: Results on SynAPT benchmark.** <sup>†</sup> Results reported in [33].

Method	Pre-training			Downstream Tasks					
	Dataset	#data	labels	UCF101	HMDB51	MiniSSV2	Diving48	IkeaFA	UAV-H
TimeSformer <sup>†</sup>	IN-21k +Synthetic	150k	✓	89.0	54.4	51.1	44.9	63.6	25.9
PPMA [74]	NH-Kinetics +Synthetic	300k	✓	<b>92.5</b>	<b>71.2</b>	67.8	64.0	<b>67.9</b>	38.5
Ours	no extra data	-		87.3	55.9	<b>69.0</b>	68.3	61.4	36.8
Ours	Shaders1k	100k		89.4	59.7	68.3	<b>72.3</b>	60.7	<b>40.0</b>

**Table 12: Results on K400.** <sup>†</sup> Results from [56]. We use 100k images from Shaders1k.

Method	Pre-training data Data	Kinetics400	
		Acc@1	Acc@5
from scratch <sup>†</sup>	-	68.8	-
VideoMAE <sup>†</sup>	K400	<b>81.5</b>	<b>95.1</b>
Ours	frames from K400	74.8	92.0
	Shaders1k	74.7	91.9

**Table 13: Comparison of accuracy on HMDB51 and UCF101 when using subsets grouped by frame difference.**


Frame difference	HMDB51	UCF101
(i) Large	32.9	68.7
(ii) Medium	<b>33.5</b>	<b>71.0</b>
(iii) Small	32.2	69.3

ing48, IkeaFA, and UAV-H. Tab. 11 presents the results. Using only synthetic images, our proposed framework partially surpasses some of the results of existing works utilizing real videos and action labels. Our framework is inferior to PPMA on UCF101, HMDB51, and IkeaFA. This is because these datasets have less data than others. PPMA leverages the 150 action labels in the video datasets for pre-training, therefore having the advantage of learning action features from a small number of videos during fine-tuning. On the other hand, our framework, not having these labels beforehand, struggled to learn meaningful features with fewer labeled data. However, our framework shows better performance on less biased datasets like MiniSSV2, Diving48, and UAV-H. This suggests that scene and object biases are mitigated when using our generated synthetic videos.

**Results on K400.** Tab. 12 shows the comparison of our framework with VideoMAE on K400. Although our framework outperforms the model 'from scratch', it falls short of the performance of VideoMAE with real videos. This shortfall is attributed to the limited diversity of pseudo-motion videos generated by PMG, especially when compared to the vast variety found in large-scale datasets. We understand our shortcoming here, but increasing the diversity of generated videos may close this gap.

#### 4.6 Performance When the Number Data for Fine-tuning is Limited

In previous experiments, the full set of video datasets for fine-tuning was available. Under these conditions, pre-training with all the videos for fine-tuning yielded better performance than our framework. However, for video datasets, there is often a limited amount of training samples to fine-tune with. To assess the effectiveness of our framework in such cases, we sampled  $\{1, 5, 10, 25, 50\}$  videos per category from HMDB51 and UCF101, respectively, and compared



**Fig. 4: Performance when the number of video data for finetuning is limited.**

the performance of our framework with VideoMAE using real videos. Fig. 4 presents the results. The model pre-trained by our framework shows higher performance compared to the model pre-trained by VideoMAE using real data. This underscores the efficacy of our framework where the available data is limited.

#### 4.7 What Does VideoMAE Learn from Pre-training with Videos?

Finally, to support our hypothesis that VideoMAE learns the correspondence of patches between frames, we conducted a simple experiment. Here, we assume that a larger frame difference in a video makes it difficult to capture this correspondence, for instance, due to extreme camera motion. Based on this, we made three subsets from HMDB51 and UCF101 depending on the frame difference; (i) videos having the top 50% average frame difference (ii) videos ranging from the 25th to the 75th percentile in average frame difference, (iii) videos having the bottom 50% average frame difference. We then use each of these subsets for pre-training, then fine-tune on the full set. The results are shown in Tab. 13. Models that are pre-trained on (i) and (iii) performed worse than those pre-trained on (ii). This lends support to our hypothesis regarding what VideoMAE learns.

## 5 Conclusion

In this paper, we introduced a self-supervised framework for pre-training video transformers solely with synthetic images. Our framework eliminates the costs associated with collecting video data and addresses concerns related to privacy, licensing, and biases inherent in real data. Our experiments have demonstrated that our framework not only outperforms existing pre-training methods with static images but also partially outperforms existing works with synthetic videos. Further analysis unveiled segments of what masked autoencoders learn from videos.

**Limitations** Our framework is inferior to pre-training with large-scale datasets like K400. We consider this to be due to a lack of fine-grained motion patterns compared to real videos. Our framework depends on hand-crafted image transformations and applies them to images globally, pseudo-motion videos do not have flexible motion patterns. Additionally, our framework does not learn high-level semantic features, because we utilized VideoMAE’s focus on capturing low-level features. Therefore, it is challenging to extend our framework to other tasks like video-text retrieval, without additional training or extra labeled data.

## References

1. Abu-El-Haija, S., Kothari, N., Lee, J., Natsev, P., Toderici, G., Varadarajan, B., Vijayanarasimhan, S.: Youtube-8m: A large-scale video classification benchmark. arXiv preprint arXiv:1609.08675 (2016)
2. Ahsan, U., Madhok, R., Essa, I.: Video jigsaw: Unsupervised learning of spatiotemporal context for video action recognition. In: 2019 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 179–189. IEEE (2019)
3. Arnab, A., Dehghani, M., Heigold, G., Sun, C., Lučić, M., Schmid, C.: Vivit: A video vision transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6836–6846 (2021)
4. Asano, Y.M., Rupprecht, C., Zisserman, A., Vedaldi, A.: Pass: An imagenet replacement for self-supervised pretraining without humans. arXiv preprint arXiv:2109.13228 (2021)
5. Bandara, W.G.C., Patel, N., Gholami, A., Nikkhah, M., Agrawal, M., Patel, V.M.: Adamae: Adaptive masking for efficient spatiotemporal learning with masked autoencoders. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14507–14517 (2023)
6. Baradad, M., Chen, R., Wulff, J., Wang, T., Feris, R., Torralba, A., Isola, P.: Procedural image programs for representation learning. *Advances in Neural Information Processing Systems* **35**, 6450–6462 (2022)
7. Baradad Jurjo, M., Wulff, J., Wang, T., Isola, P., Torralba, A.: Learning to see by looking at noise. *Advances in Neural Information Processing Systems* **34**, 2556–2569 (2021)
8. Benaim, S., Ephrat, A., Lang, O., Mosseri, I., Freeman, W.T., Rubinstein, M., Irani, M., Dekel, T.: Speednet: Learning the speediness in videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9922–9931 (2020)
9. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: ICML. vol. 2, p. 4 (2021)
10. Buolamwini, J., Gebru, T.: Gender shades: Intersectional accuracy disparities in commercial gender classification. In: Conference on fairness, accountability and transparency. pp. 77–91. PMLR (2018)
11. Caba Heilbron, F., Escorcia, V., Ghanem, B., Carlos Nibbles, J.: Activitynet: A large-scale video benchmark for human activity understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 961–970 (2015)
12. Chen, X., Xie, S., He, K.: An empirical study of training self-supervised vision transformers. arXiv preprint arXiv:2104.02057 **2**(5), 6 (2021)
13. Choi, J., Gao, C., Messou, J.C., Huang, J.B.: Why can’t i dance in the mall? learning to mitigate scene bias in action recognition. *Advances in Neural Information Processing Systems* **32** (2019)
14. De Souza, C., Gaidon, A., Cabon, Y., López, A.M.: Procedural generation of videos to train deep action recognition networks. corr. arXiv preprint arXiv:1612.00881 (2016)
15. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
16. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)

17. Feichtenhofer, C., Fan, H., Li, Y., He, K.: Masked autoencoders as spatiotemporal learners. arXiv preprint arXiv:2205.09113 (2022)
18. Feichtenhofer, C., Fan, H., Xiong, B., Girshick, R., He, K.: A large-scale study on unsupervised spatiotemporal representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3299–3309 (2021)
19. Fernando, B., Bilen, H., Gavves, E., Gould, S.: Self-supervised video representation learning with odd-one-out networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3636–3645 (2017)
20. Fischer, P., Dosovitskiy, A., Ilg, E., Häusser, P., Hazırbaş, C., Golkov, V., Van der Smagt, P., Cremers, D., Brox, T.: FlowNet: Learning optical flow with convolutional networks. arXiv preprint arXiv:1504.06852 (2015)
21. Ghadiyaram, D., Tran, D., Mahajan, D.: Large-scale weakly-supervised pre-training for video action recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12046–12055 (2019)
22. Goyal, R., Ebrahimi Kahou, S., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haenel, V., Fruend, I., Yianilos, P., Mueller-Freitag, M., et al.: The "something something" video database for learning and evaluating visual common sense. In: Proceedings of the IEEE international conference on computer vision. pp. 5842–5850 (2017)
23. Guo, X., Wu, W., Wang, D., Su, J., Su, H., Gan, W., Huang, J., Yang, Q.: Learning video representations of human motion from synthetic data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20197–20207 (2022)
24. Hara, K., Kataoka, H., Satoh, Y.: Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 6546–6555 (2018)
25. Huang, Z., Zhang, S., Jiang, J., Tang, M., Jin, R., Ang, M.H.: Self-supervised motion learning from static images. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1276–1285 (2021)
26. Hwang, H., Jang, C., Park, G., Cho, J., Kim, I.J.: Eldersim: A synthetic data generation platform for human action recognition in eldercare applications. IEEE Access (2021)
27. Jenni, S., Meishvili, G., Favaro, P.: Video representation learning by recognizing temporal transformations. In: European Conference on Computer Vision. pp. 425–442. Springer (2020)
28. Kataoka, H., Hara, K., Hayashi, R., Yamagata, E., Inoue, N.: Spatiotemporal initialization for 3d cnns with generated motion patterns. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1279–1288 (2022)
29. Kataoka, H., Hayamizu, R., Yamada, R., Nakashima, K., Takashima, S., Zhang, X., Martinez-Noriega, E.J., Inoue, N., Yokota, R.: Replacing labeled real-image datasets with auto-generated contours. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21232–21241 (2022)
30. Kataoka, H., Okayasu, K., Matsumoto, A., Yamagata, E., Yamada, R., Inoue, N., Nakamura, A., Satoh, Y.: Pre-training without natural images. In: Proceedings of the Asian Conference on Computer Vision (2020)
31. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017)



32. Kim, D., Cho, D., Kweon, I.S.: Self-supervised video representation learning with space-time cubic puzzles. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 33, pp. 8545–8552 (2019)
33. Kim, Y.w., Mishra, S., Jin, S., Panda, R., Kuehne, H., Karlinsky, L., Saligrama, V., Saenko, K., Oliva, A., Feris, R.: How transferable are video representations based on synthetic data? *Advances in Neural Information Processing Systems* **35**, 35710–35723 (2022)
34. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: Hmdb: a large video database for human motion recognition. In: *2011 International conference on computer vision*. pp. 2556–2563. IEEE (2011)
35. Li, K., Wang, Y., Li, Y., Wang, Y., He, Y., Wang, L., Qiao, Y.: Unmasked teacher: Towards training-efficient video foundation models. *arXiv preprint arXiv:2303.16058* (2023)
36. Li, T., Liu, J., Zhang, W., Ni, Y., Wang, W., Li, Z.: Uav-human: A large benchmark for human behavior understanding with unmanned aerial vehicles. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 16266–16275 (2021)
37. Li, Y., Vasconcelos, N.: Repair: Removing representation bias by dataset resampling. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9572–9581 (2019)
38. Li, Y., Li, Y., Vasconcelos, N.: Resound: Towards action recognition without representation bias. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 513–528 (2018)
39. Ma, Z., Zhang, Z., Chen, Y., Qi, Z., Luo, Y., Li, Z., Yuan, C., Li, B., Qie, X., Shan, Y., et al.: Order-prompted tag sequence generation for video tagging. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 15681–15690 (2023)
40. Mahajan, D., Girshick, R., Ramanathan, V., He, K., Paluri, M., Li, Y., Bharambe, A., Van Der Maaten, L.: Exploring the limits of weakly supervised pretraining. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 181–196 (2018)
41. Miech, A., Zhukov, D., Alayrac, J.B., Tapaswi, M., Laptev, I., Sivic, J.: Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2630–2640 (2019)
42. Mu, J., Qiu, W., Hager, G.D., Yuille, A.L.: Learning from synthetic animals. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12386–12395 (2020)
43. Nakamura, R., Kataoka, H., Takashima, S., Noriega, E.J.M., Yokota, R., Inoue, N.: Pre-training vision transformers with very limited synthesized images. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 20360–20369 (2023)
44. Nakashima, K., Kataoka, H., Matsumoto, A., Iwata, K., Inoue, N., Satoh, Y.: Can vision transformers learn without natural images? In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 36, pp. 1990–1998 (2022)
45. Pan, T., Song, Y., Yang, T., Jiang, W., Liu, W.: Videomoco: Contrastive video representation learning with temporally adversarial examples. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11205–11214 (2021)
46. Perlin, K.: An image synthesizer. *ACM Siggraph Computer Graphics* **19**(3), 287–296 (1985)

47. Perlin, K.: Improving noise. In: Proceedings of the 29th annual conference on Computer graphics and interactive techniques. pp. 681–682 (2002)
48. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems* **35**, 25278–25294 (2022)
49. Shu, F., Chen, B., Liao, Y., Xiao, S., Sun, W., Li, X., Zhu, Y., Wang, J., Liu, S.: Masked contrastive pre-training for efficient video-text retrieval. *arXiv preprint arXiv:2212.00986* (2022)
50. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402* (2012)
51. Roberto de Souza, C., Gaidon, A., Cabon, Y., Manuel Lopez, A.: Procedural generation of videos to train deep action recognition networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4757–4767 (2017)
52. Sun, C., Shrivastava, A., Singh, S., Gupta, A.: Revisiting unreasonable effectiveness of data in deep learning era. In: Proceedings of the IEEE international conference on computer vision. pp. 843–852 (2017)
53. Sun, X., Chen, P., Chen, L., Li, C., Li, T.H., Tan, M., Gan, C.: Masked motion encoding for self-supervised video representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2235–2245 (2023)
54. Takashima, S., Hayamizu, R., Inoue, N., Kataoka, H., Yokota, R.: Visual atoms: Pre-training vision transformers with sinusoidal waves. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18579–18588 (2023)
55. Teed, Z., Deng, J.: 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. pp. 402–419. Springer (2020)
56. Tong, Z., Song, Y., Wang, J., Wang, L.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602* (2022)
57. Toyer, S., Cherian, A., Han, T., Gould, S.: Human pose forecasting via deep markov models. In: 2017 International Conference on Digital Image Computing: Techniques and Applications (DICTA). pp. 1–8. IEEE (2017)
58. Varol, G., Laptev, I., Schmid, C., Zisserman, A.: Synthetic humans for action recognition from unseen viewpoints. *International Journal of Computer Vision* **129**(7), 2264–2287 (2021)
59. Wang, J., Jiao, J., Liu, Y.H.: Self-supervised video representation learning by pace prediction. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16. pp. 504–521. Springer (2020)
60. Wang, L., Huang, B., Zhao, Z., Tong, Z., He, Y., Wang, Y., Wang, Y., Qiao, Y.: Videomae v2: Scaling video masked autoencoders with dual masking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14549–14560 (2023)
61. Wang, R., Chen, D., Wu, Z., Chen, Y., Dai, X., Liu, M., Jiang, Y.G., Zhou, L., Yuan, L.: Bevt: Bert pretraining of video transformers. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14733–14743 (2022)

62. Wang, R., Chen, D., Wu, Z., Chen, Y., Dai, X., Liu, M., Yuan, L., Jiang, Y.G.: Masked video distillation: Rethinking masked feature modeling for self-supervised video representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6312–6322 (2023)
63. Wang, T., Zhao, J., Yatskar, M., Chang, K.W., Ordonez, V.: Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5310–5319 (2019)
64. Wei, C., Fan, H., Xie, S., Wu, C.Y., Yuille, A., Feichtenhofer, C.: Masked feature prediction for self-supervised visual pre-training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14668–14678 (2022)
65. Wu, Z., Wang, Z., Wang, Z., Jin, H.: Towards privacy-preserving visual recognition via adversarial training: A pilot study. In: Proceedings of the European conference on computer vision (ECCV). pp. 606–624 (2018)
66. Xu, D., Xiao, J., Zhao, Z., Shao, J., Xie, D., Zhuang, Y.: Self-supervised spatiotemporal learning via video clip order prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10334–10343 (2019)
67. Yang, H., Huang, D., Wen, B., Wu, J., Yao, H., Jiang, Y., Zhu, X., Yuan, Z.: Self-supervised video representation learning with motion-aware masked autoencoders. arXiv preprint arXiv:2210.04154 (2022)
68. Yang, K., Qinami, K., Fei-Fei, L., Deng, J., Russakovsky, O.: Towards fairer datasets: Filtering and balancing the distribution of the people subtree in the imagenet hierarchy. In: Proceedings of the 2020 conference on fairness, accountability, and transparency. pp. 547–558 (2020)
69. Yao, Y., Liu, C., Luo, D., Zhou, Y., Ye, Q.: Video playback rate perception for self-supervised spatio-temporal representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6548–6557 (2020)
70. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6023–6032 (2019)
71. Yun, S., Oh, S.J., Heo, B., Han, D., Kim, J.: Videomix: Rethinking data augmentation for video classification. arXiv preprint arXiv:2012.03457 (2020)
72. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. arXiv preprint arXiv:1710.09412 (2017)
73. Zhang, Z., Chen, Y., Ma, Z., Qi, Z., Yuan, C., Li, B., Shan, Y., Hu, W.: Create: A benchmark for chinese short video retrieval and title generation. arXiv preprint arXiv:2203.16763 (2022)
74. Zhong, H., Mishra, S., Kim, D., Jin, S., Panda, R., Kuehne, H., Karlinsky, L., Saligrama, V., Oliva, A., Feris, R.: Learning human action recognition representations without real humans. arXiv preprint arXiv:2311.06231 (2023)