

VideoMAC: Video Masked Autoencoders Meet ConvNets

Gensheng Pei*, Tao Chen*, Xiruo Jiang, Huafeng Liu, Zeren Sun,† Yazhou Yao†

School of Computer Science and Engineering, Nanjing University of Science and Technology

<https://github.com/NUST-Machine-Intelligence-Laboratory/VideoMAC>

Abstract

Recently, the advancement of self-supervised learning techniques, like masked autoencoders (MAE), has greatly influenced visual representation learning for images and videos. Nevertheless, it is worth noting that the predominant approaches in existing masked image / video modeling rely excessively on resource-intensive vision transformers (ViTs) as the feature encoder. In this paper, we propose a new approach termed as **VideoMAC**, which combines video masked autoencoders with resource-friendly ConvNets. Specifically, VideoMAC employs symmetric masking on randomly sampled pairs of video frames. To prevent the issue of mask pattern dissipation, we utilize ConvNets which are implemented with sparse convolutional operators as encoders. Simultaneously, we present a simple yet effective masked video modeling (MVM) approach, a dual encoder architecture comprising an online encoder and an exponential moving average target encoder, aimed to facilitate inter-frame reconstruction consistency in videos. Additionally, we demonstrate that VideoMAC, empowering classical (ResNet) / modern (ConvNeXt) convolutional encoders to harness the benefits of MVM, outperforms ViT-based approaches on downstream tasks, including video object segmentation (+5.2% / 6.4% $J&F$), body part propagation (+6.3% / 3.1% mIoU), and human pose tracking (+10.2% / 11.1% PCK@0.1).

1. Introduction

Spurred by the success of masked language modeling in natural language processing [3, 10, 12], transformer-based visual models, notably vision transformers (ViT) [14] adopt masked image modeling (MIM) are attracting attention [1, 2, 15, 23, 34, 67]. This strategy involves masking or corrupting a portion of information within an image and then requiring the model to recover the masked portion, encouraging the model to learn useful image representations.

*Equal contribution.

†Corresponding author.

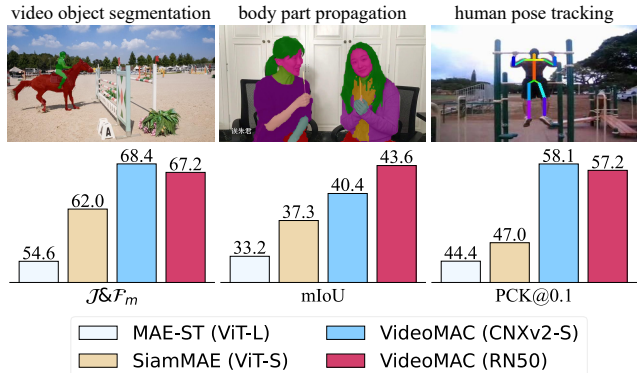


Figure 1. State-of-the-art MAE architectures (e.g., MAE-ST [16] and SiamMAE [20]) in masked video modeling commonly employ ViT-based encoders. We propose **VideoMAC**, a new video masked autoencoder built using pure ConvNets. In this study, we model VideoMAC with two of the most emblematic families of ConvNets, namely ConvNeXtv2 [64] and ResNet [21]. Notably, VideoMAC exhibits superior performance on a range of downstream tasks, e.g., video object segmentation on DAVIS17 [46], body part propagation on VIP [73], and human pose tracking on JHMDB [26], compared to ViT-based methods [16, 20].

With the successful application of MIM, some pioneering works [16, 56] have shifted the focus of video representation learning to masked video modeling (MVM), yielding promising results in downstream video tasks. However, features learned by MVM are designed explicitly for pixel reconstruction tasks. They perform well when fine-tuned for downstream video tasks but exhibit subpar performance without specific fine-tuning [20]. Their performance is even inferior to MIM methods [23, 30], and it lags significantly behind supervised methods [21, 36].

We strive to discern the factors contributing to this issue. To the best of our knowledge, most current MVM methods are based on the isotropic design of ViT [16, 20, 56, 65]. The reason is that early MIM / MVM methods are explicitly developed for ViT, as they allow the use of only visible patches as the sole input to the encoder, which aligns well with the ViT architecture. For instance, the representative approach, masked autoencoders (MAE) [23], enhances the efficiency and effectiveness of model pre-training by focusing predominantly on the unmasked regions. However,




Figure 2. Visualization of the heatmap during the reconstruction of masked patches in a random frame pair. It is evident that our approach highlights similar regions for both (a) past and (b) current frames, proficiently reconstructing colors and contours.

previous MVM approaches [16, 56], confined by the plain vision transformer design devoid of a hierarchical structure tend to result in inadequate spatial modeling capabilities for local features [17, 35, 49, 62, 70]. Such capabilities are crucial for dense downstream tasks like segmentation and detection. Moreover, the uniform reliance on ViT in existing MVM methods has resulted in significant computational resource consumption [16, 56, 58]. One might first consider ConvNets [21, 28, 29, 36, 50, 53], characterized by their low resource consumption and built-in hierarchical structure, as an ideal solution to tackle the aforementioned issues. Nevertheless, ConvNets appears understudied in MVM. Consequently, it raises the question: *can ConvNets, endowed with a pyramid vision architecture, guide MVM towards the era of hierarchical pre-training?*

To address this concern, we propose a simple yet potent MVM approach, referred to as *Video Masked Autoencoders with ConvNets (VideoMAC)*. Uniquely, our method represents an initial endeavor to supplement video pre-training by integrating hierarchical ConvNets. However, applying ConvNets to MVM proves more complex than utilizing ViTs. Typical ConvNets might be incompatible with MVM due to the employment of dense sliding windows with overlap, resulting in information leakage within masked regions. To efficiently incorporate mask modeling into ConvNets, we opt for sparse convolution over dense convolution to restructure the original ConvNets, an approach validated in recent studies [54, 64] to effectively maintain the structure of the masked region when utilizing ConvNets. Another hurdle for MVM is the effective incorporation of temporal data. Current MVM pre-training methods [16, 56] involving multi-frame inputs are primarily used in downstream resource-intensive tasks such as video action recognition. Although methods using two-frame inputs for pre-training lower resource usage [20, 65], they typically employ the Siamese architecture, necessitating robust skills for designing reconstruction targets. Importantly, all aforementioned methods utilize computationally intensive, non-hierarchical ViT [14] as their encoders. To this end, we utilize pairs of frames as inputs and introduce the reconstruction consis-

tency loss, a streamlined extension of MAE to accommodate spatio-temporal data. We formulate online networks through gradient optimization and establish target networks by updating parameters via an exponential moving average (EMA). This allows for individual recovery of masked frame pairs and computation of the reconstruction consistency loss between them. As illustrated in Fig. 2, our approach can focus on similar regions between two frames and reconstruct the masked portions efficiently.

Notably, our approach transcends previous state-of-the-art techniques in three downstream tasks, *e.g.*, video object segmentation on DAVIS17 [46], body part propagation on VIP [73], and human pose tracking on JHMDB [26]. Moreover, we present promising results utilizing VideoMAC in fine-tuning for image recognition on ImageNet1K [11]. To summarize, our main contributions of this work include:

- We propose a simple yet potent solution of masked video autoencoders, marking the initial endeavor to successfully implement MVM using pure ConvNets.
- We introduce the reconstruction consistency loss to enable elegant modeling of spatio-temporal data.
- Our approach significantly outperforms existing ViT-based MVM methods in various downstream tasks.

While the focus of the computer vision community is increasingly shifting towards vision transformers, we trust that VideoMAC will rekindle interest in exploring ConvNet-based MVM, thereby catalyzing the discovery of new potentials across a variety of video tasks.

2. Related Work

Masked Image Modeling. Self-supervised learning (SSL) has been rapidly developing in computer vision, demonstrating substantial potential for deriving valuable representations from extensive, unlabeled datasets. A key advantage of SSL [2, 5–8, 18, 22, 48, 57, 69, 72] is its independence from costly human annotations, thereby providing robust prior knowledge across various visual tasks. The success of masked language modeling in NLP [3, 10, 12, 13, 47] has catalyzed a rising interest in masked image modeling (MIM) for visual pre-training. This strategy involves masking or corrupting a portion of information within an image and then requiring the model to recover the masked portion, encouraging the model to learn useful image representations [2, 23, 67]. Notably, pioneering works [54, 64] have recently applied ConvNets to MIM pre-training, paving the way for masking modeling methods based on ConvNets.

Masked Video Modeling. Recently, researchers have embarked on exploring MVM approaches [16, 19, 65, 66] to comprehend unsupervised video representations, exhibiting effectiveness across a multitude of downstream tasks [4, 25, 31, 32, 40–44, 60, 61, 71, 74, 75]. However, current MVM methodologies [20, 56, 58, 59] primarily depend on the isotropic ViT, a visual backbone compatible with mask-

ing modeling. Furthermore, these ViT-based MVM strategies may encounter challenges with inadequate local feature modeling when dealing with dense tasks such as segmentation and detection due to their non-hierarchical structure [35, 37, 62, 70]. To the best of our knowledge, no MVM approaches are explicitly designed for hierarchical network-based architectures, especially ConvNets.

It is worth noting that while the isotropic ViT is well suited for the implementation of MVM methods [16, 19, 56], this is exceptionally difficult for ConvNets. In this paper, our VideoMAC is committed to developing an efficient video pre-training tool that can equip ConvNets with superior, or at least comparable, performance to ViTs.

3. Approach

Our VideoMAC builds upon MAE [23] and, pioneeringly, facilitates masked video modeling with a purely convolutional architecture. As depicted in Fig. 3, VideoMAC is designed to reconstruct masked patches, which are randomly and symmetrically masked on frame pairs at a high masking ratio (0.75 as used in this study). The framework achieves such reconstructions by exclusively delivering visible patches to both the online and target encoders. In this section, we describe its main components in detail.

3.1. Preliminaries

Recently, masked autoencoders (MAE) [23] introduces a novel methodology to the domain of self-supervised visual representation learning. The primary process involves training a model to reconstruct masked areas in images whose contents are partially occluded. The fundamental architecture of a standard MAE model is composed of three core elements: the masking design, an encoder, and a decoder. In mathematical terms, given an image $\mathbf{I} \in \mathbb{R}^{3 \times H \times W}$, the MAE model patchifies \mathbf{I} into N distinct, non-overlapping patches, represented by $\hat{\mathbf{I}} \in \mathbb{R}^{N \times (P^2 \times 3)}$, each with a size of P . Masks are randomly applied to a subset of these patches, denoted by the subset \mathcal{M} .

The encoder $f_{\mathcal{E}}$ processes the visible patches \mathcal{V} , producing latent representations that are represented as $z = f_{\mathcal{E}}(\mathcal{V})$. Then, the decoder $f_{\mathcal{D}}$ attempts to restore the masked subset \mathcal{M} utilizing these latent representations, generating $\mathbf{O} = f_{\mathcal{D}}(z; \mathcal{M})$, where $\mathbf{O} \in \mathbb{R}^{N \times (P^2 \times 3)}$ depicts the reconstructed patches. The MAE model utilizes the mean-squared error (MSE) to compute the reconstruction loss \mathcal{L}_{rec} , which is solely calculated based on the masked patches and is formulated as $\mathcal{L}_{rec} = \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \|\hat{\mathbf{I}}_i - \mathbf{O}_i\|_2^2$.

3.2. Video Masked Autoencoder with ConvNets

Masking. In this work, we randomly select two consecutive frames, $\mathbf{I}^{(1)}, \mathbf{I}^{(2)} \in \mathbb{R}^{3 \times H \times W}$, and divide them into N non-overlapping patches, represented as $\hat{\mathbf{I}}^{(1)}, \hat{\mathbf{I}}^{(2)} \in \mathbb{R}^{N \times (P^2 \times 3)}$.

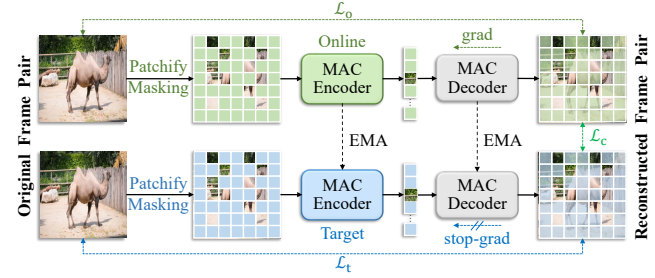


Figure 3. An illustration of VideoMAC for ConvNet-based MVM. During pre-training, we mask 75% of symmetric patches from two frames randomly. In our VideoMAC, the MVM of frame pairs is achieved by an online network optimized by gradients (■, online loss \mathcal{L}_o) and a target network updated by EMA (■, target loss \mathcal{L}_t). \mathcal{L}_c is computed as the reconstruction consistency loss between reconstructed patches of frame pairs.

Regarding the mask configuration, our approach adopts symmetric masking (Table 3d) to guarantee consistency in the positions of the masked areas between frame pairs. Given the consecutive frames, $\mathcal{M}^{(1)}$ and $\mathcal{M}^{(2)}$ constitute the subsets of masked patches, whereas $\mathcal{V}^{(1)}$ and $\mathcal{V}^{(2)}$ correspond to the visible patches.

It is inherent that images and videos possess a greater level of information redundancy [16, 22] compared to languages. Accordingly, we employ a masking ratio of 0.75 (Table 4f) in our methodology. Subsequent empirical analyses (Sec. 4.2) have demonstrated an improvement in performance with higher masking ratios. Nevertheless, employing excessively high masking ratios may induce the risk of training collapse, thereby leading to sudden and substantial decreases in model performance.

MAC Encoder. In ViT-based models, the prevention of information leakage from masked regions is comparatively straightforward [64, 70]. These models solely employ the visible patches as the exclusive input to the encoder, thereby allowing ViT-based approaches [16, 20, 56, 65] to amplify the efficiency and efficacy of model pre-training by concentrating dedicatedly on unmasked regions during this phase. In contrast, achieving a similar task with ConvNets poses a more intricate dilemma, necessitating the maintenance of the 2D image structure. There are two naive solutions: (i) the instigation of learnable masked markers at the input stage, and (ii) the transformation of the hierarchical convolutional network into an isotropic structure [36] that mirrors the ViT-style technique (Table 2). Regardless, both strategies may decrease pre-training efficiency and possibly engender inconsistencies between training and testing.

This paper is at the forefront of proposing an entirely ConvNet-based MVM approach to address this challenge. Drawing inspiration from [54, 64], we utilize sparse convolutional networks to develop a masked autoencoder, which usurps traditional dense convolution. This strategy aims to meet three primary goals: (i) maintaining the structural integrity of the image during feature encoding, (ii) preventing

mask dissipation in the course of hierarchical convolution, and (iii) markedly amplifying the computational efficiency of MVM. Specifically, our VideoMAC uses a fully convolutional architecture. Conventionally, ConvNets are structured with four stages, where the resolution of features is halved after each stage, representing a patch size of 32. For an input frame size of 224×224 , the features are sequentially downsampled at every stage to achieve resolutions of 56×56 , 28×28 , 14×14 , and 7×7 , respectively. Subsequently, masks are produced and recursively upsampled to the highest resolution (*i.e.*, 56×56) in the final stage.

Current MVMs typically necessitate significant computational resources, posing a challenge for practical applications. Our strategy, which substitutes ConvNets for ViTs, does mitigate computational complexity. Furthermore, studying existing MVM methods, we observe that to capture temporal signals within videos, most methods either utilize multi-frame data input or opt for a siamese encoder structure. We diverge from these widespread strategies and instead adopt a more resource-efficient method, namely the online-target encoder architecture (Fig. 3), commonly applied in self-supervised image representation learning. The target encoder’s weights are updated by the online encoder via EMA. This architectural integration into our MVM encoding process results in three-fold advantages: (i) considerable reduction of computational complexity, (ii) joint optimization of frame-to-feature representation by the online and target encoders, and (iii) simplicity in incorporating existing ConvNets into our framework, thus improving the performance of downstream tasks. Consequently, visible patches $\mathcal{V}^{(1)}$ and $\mathcal{V}^{(2)}$ of consecutive frames are processed by the online encoder $f_{\mathcal{E}}^o$ and the target encoder $f_{\mathcal{E}}^t$ to generate latent representations $z^{(1)} = f_{\mathcal{E}}^o(\mathcal{V}^{(1)})$ and $z^{(2)} = f_{\mathcal{E}}^t(\mathcal{V}^{(2)})$, respectively.

MAC Decoder. Unlike previous methods that utilize hierarchical decoders [54, 55] or transformer decoders [16, 33, 56], our VideoMAC implements a lightweight decoder derived from the ConvNeXt block [36]. This simplifies the design and bolsters the efficiency of the pre-training process. Analogous to our encoder, we construct an online decoder $f_{\mathcal{D}}^o$ and a target decoder $f_{\mathcal{D}}^t$, further employing EMA for parameter updates (Fig. 3). Given the latent representations $z^{(1)}$ and $z^{(2)}$ from two successive frames, decoders $f_{\mathcal{D}}^o$ and $f_{\mathcal{D}}^t$ reconstruct masked patches $\mathcal{M}^{(1)}$ and $\mathcal{M}^{(2)}$ to yield $\mathcal{O}^{(1)} = f_{\mathcal{D}}^o(z^{(1)}; \mathcal{M}^{(1)})$ and $\mathcal{O}^{(2)} = f_{\mathcal{D}}^t(z^{(2)}; \mathcal{M}^{(2)})$, respectively. In this paper, we employ a single-block design for the decoder, with the decoder feature projection dimension set to 512. Detailed ablation analysis can be found in Sec. 4.2, and Tables 3b and 3c.

Reconstruction Consistency. In line with the majority of MAE-based methods, we compute the Mean Squared Error (MSE) between the reconstituted and original images to assess the reconstruction loss. We elect to aim for each patch

of normalized pixels, executing the loss calculation solely for masked patches. As a result, we synchronize the values of normalized pixels within each image patch, applying the loss function exclusively to those patches subject to masking. Significantly, our VideoMAC yields two frames of reconstructed results from both the online and target encoders, necessitating individual loss computations:

$$\begin{aligned} \mathcal{L}_o &= \frac{1}{|\mathcal{M}^{(1)}|} \sum_{i \in \mathcal{M}^{(1)}} \|\hat{I}_i^{(1)} - \mathcal{O}_i^{(1)}\|_2^2, \\ \mathcal{L}_t &= \frac{1}{|\mathcal{M}^{(2)}|} \sum_{i \in \mathcal{M}^{(2)}} \|\hat{I}_i^{(2)} - \mathcal{O}_i^{(2)}\|_2^2, \end{aligned} \quad (1)$$

where \mathcal{L}_o and \mathcal{L}_t denote online and target losses correspondingly. Due to the correlation between consecutive video frames, we pursue temporal consistency in addition to the reconstruction loss. Here, we introduce an intuitive assumption: if the pixel deviation between the two raw frames tends to a particular value, denoted as ϵ , it is inferred that the divergence between the reconstructed frames should also trend towards ϵ . Based on this assumption, we can estimate the reconstruction consistency between two frames and use it as a proxy for temporal consistency. Hence, the reconstruction consistency loss \mathcal{L}_c is thus expressed as:

$$\mathcal{L}_c = \frac{1}{|\mathcal{M}^{(1)}|} \sum_{i \in \mathcal{M}^{(1)}, j \in \mathcal{M}^{(2)}} \|\mathcal{O}_i^{(1)} - \mathcal{O}_j^{(2)}\|_2^2, \quad (2)$$

where $|\mathcal{M}^{(1)}|$ can be substituted by $|\mathcal{M}^{(2)}|$, courtesy of our utilization of symmetric masking. Ultimately, our approach culminates in a total loss expressed as:

$$\mathcal{L}_{total} = \mathcal{L}_o + \mathcal{L}_t + \gamma \mathcal{L}_c, \quad (3)$$

with γ symbolizing the weight factor of the consistency loss (see Table 4e and Fig. 7 for detailed analysis). Accordingly, our approach, VideoMAC, accomplishes the realization of a fully convolutional network MVM. Moreover, the reconstruction consistency loss introduced in this paper provides a promising solution for modeling spatio-temporal data.

3.3. Implementation

Architecture. Our VideoMAC possesses the adaptability to employ any convolutional network the an encoder. To permit a fair comparison, we choose two representative convolutional networks for our experiments: ResNet (RN) [21] and ConvNeXtv2 (CNXv2) [64]. Although RN50 (26M) and CNXv2-S (50M) serve as the primary experimental models, we further investigate models of diverse sizes (Tables 2, 3a and 5). In this paper, the decoder component involves a readily available ConvNeXt block [36], thereby dispensing with the necessity for intricate network design. During pre-training, we employ an encoder implemented

Method	Type	Backbone	Training dataset	Epoch	DAVIS17 [46]			VIP [73]	JHMDB [26]	
					$\mathcal{J}\&\mathcal{F}_m$	\mathcal{J}_m	\mathcal{F}_m	mIoU	PCK@0.1	PCK@0.2
Supervised [21] _{CVPR16}	SL	RN50	IN1K (1.28M, -)	100	66.0	63.7	68.4	39.5	59.2	78.3
DenseCL [63] _{CVPR21}	CL	RN50	IN1K (1.28M, -)	200	61.1	59.8	62.6	32.5	56.1	78.0
DINO [6] _{ICCV21}	CL	ViT-S/16	IN1K (1.28M, -)	800	61.8	60.2	63.4	36.2	45.6	75.0
ODIN ² [24] _{ECCV22}	CL	RN50	IN1K (1.28M, -)	1000	54.1	54.3	53.9	-	-	-
CrOC [52] _{CVPR23}	CL	ViT-S/16	IN1K (1.28M, -)	300	44.7	43.5	45.9	26.1	-	-
MAE [23] _{CVPR22}	MIM	ViT-B/16	IN1K (1.28M, -)	1600	53.5	52.1	55.0	28.1	44.6	73.4
FCMAE [64] _{CVPR23}	MIM	CNXv2-B	IN1K (1.28M, -)	1600	43.7	41.9	45.5	24.9	51.7	74.2
RC-MAE [30] _{ICLR23}	MIM	ViT-S/16	IN1K (1.28M, -)	1600	49.2	48.9	50.5	29.7	43.2	72.3
SparK [54] _{ICLR23}	MIM	RN50	IN1K (1.28M, -)	1600	55.6	56.3	54.9	32.8	52.9	75.2
MAE-ST [16] _{NeurIPS22}	MVM	ViT-L/16	K400 (-, 800 hours)	1600	54.6	55.5	53.6	33.2	44.4	72.5
VideoMAE [56] _{NeurIPS22}	MVM	ViT-S/16	K400 (-, 800 hours)	1600	39.3	39.7	38.9	22.3	41.0	67.9
DropMAE [65] _{CVPR23}	MVM	ViT-B/16	K400 (-, 800 hours)	1600	53.4	51.8	55.0	31.1	42.3	69.2
SiamMAE [20] _{NeurIPS23}	MVM	ViT-S/16	K400 (-, 800 hours)	2000	62.0	60.3	63.7	37.3	47.0	76.1
<i>Ours</i>										
VideoMAC	MVM	RN50	YT18 (-, 5.58 hours)	100	<u>67.2</u>	<u>64.9</u>	<u>69.5</u>	43.6	57.2	<u>79.2</u>
VideoMAC	MVM	CNXv2-S	YT18 (-, 5.58 hours)	100	68.4	65.3	71.4	<u>40.4</u>	<u>58.1</u>	80.0

SL: Supervised Learning. CL: Contrastive Learning. IN1K: ImageNet1K [11]. K400: Kinetics-400 [27]. YT18: YouTube-VOS 2018 [68].

Table 1. **Comparing VideoMAC with previous supervised and self-supervised approaches (SL, CL, MIM, and MVM)** on downstream tasks, e.g., video object segmentation on DAVIS17 [46], body part propagation on VIP [73], and human pose tracking on JHMDB [26].

MVM Method	Resolution	#Param.	FLOPs	$\mathcal{J}\&\mathcal{F}_m$
<i>Iso. ViT Backbone</i>				
MAE-ST [16]	14 × 14	304M	61.6G	54.6
VideoMAE [56]	14 × 14	<u>22M</u>	4.6G	39.3
DropMAE [65]	14 × 14	87M	17.6G	53.4
SiamMAE [20]	14 × 14	<u>22M</u>	4.6G	62.0
<i>Iso. ConvNet Backbone</i>				
VideoMAC (CNXv2-S [†])	14 × 14	<u>22M</u>	4.3G	53.8
<i>Hie. ConvNet Backbone</i>				
VideoMAC (RN18)	7 × 7	12M	1.8G	64.7
VideoMAC (RN50)	7 × 7	26M	<u>4.1G</u>	67.2
VideoMAC (CNXv2-T)	7 × 7	29M	4.5G	<u>67.5</u>
VideoMAC (CNXv2-S)	7 × 7	50M	8.7G	68.4

Table 2. **Comparisons with isotropic (iso.) and hierarchical (hie.) based MVM methods.** ‘†’ denotes our modified *ios.* version.

with the MinkowskiEngine sparse convolution library [9]. It’s noteworthy that the sparse convolution layer can be seamlessly converted back to a standard convolution layer in downstream tasks without requiring additional processes.

Pre-training. Different from previous methods that pre-train on the large-scale video dataset K400 [27], VideoMAC pre-trains on YouTube-VOS [68] for 100 epochs, leading to a significant reduction in pre-training time by over 140 times. During VideoMAC pre-training, the standard input comprises two frames with a spatial size of 224×224 pixels. The frame interval for this study is determined to be 1 (Table 4c), implying consecutive frames are chosen. We employ minimal data augmentation, encompassing only random resized cropping and horizontal flipping. A masking ratio of 0.75 is utilized to pre-train our models. We conduct training using an AdamW optimizer [39] with a batch size of 512 and a learning rate of $2e-3$. Cosine decay [38] is applied to adjust the learning rate.

Downstream Tasks. To gauge the efficacy of VideoMAC at video representation learning, we undertake comprehensive experiments on three video downstream tasks, namely video object segmentation, body part propagation, and hu-

man pose tracking (Table 1). We also investigate the performance of image recognition through video pre-training. Specifically, our backbone is pre-trained on video data and then fine-tuned on ImageNet1K [11] (Table 5).

4. Experiments

This section begins with assessing VideoMAC’s performance in three downstream tasks, comparing it to previous state-of-the-art methods. Subsequently, we conduct extensive ablation studies to elucidate the effectiveness of our proposed approach. The final part of this section explores VideoMAC’s ability for image recognition, accompanied by a visualization of the reconstruction results.

4.1. Comparison with Previous methods

Video Object Segmentation. We evaluate VideoMAC on the popular DAVIS17 [46] dataset, which includes a validation set comprising 30 videos. Our evaluation methodology aligns with [6, 25], and we report the typical evaluation metrics [45], including the mean region similarity \mathcal{J}_m , the mean boundary accuracy \mathcal{F}_m , and their average values $\mathcal{J}\&\mathcal{F}_m$. As shown in Table 1, our VideoMAC achieves state-of-the-art results for video object segmentation. With a smaller video training dataset and fewer training epochs, both our VideoMAC models (*i.e.*, CNXv2-S and RN50) surpass the SiamMAE [20] by **6.4%** and **5.2%**, respectively, in terms of $\mathcal{J}\&\mathcal{F}_m$. It’s worth noting that our approach achieves superior performance to supervised method [21] based on MVM, with improvements of **2.4%** and **1.4%**, respectively. Owing to the hierarchical feature extraction capabilities inherent in ConvNets, our approach demonstrates superior proficiency in spatial information extraction in contrast to ViT-based MVM methods [16, 20, 56, 65]. Concurrently, the introduced reconstruction consistency framework strengthens the acquisition of temporal correlations.




Figure 4. **Qualitative results of our VideoMAC** (using CNXv2-S) for three video downstream tasks: (a) video object segmentation on DAVIS17 [46]), (b) body part propagation on VIP [73], and (c) human pose tracking on JHMDB [26].

Backbone	$\mathcal{J} \& \mathcal{F}_m$	\mathcal{J}_m	\mathcal{F}_m
R18	64.7	62.8	66.5
R50	67.2	64.9	69.5
CNXv2-T	67.5	64.4	70.6
CNXv2-S	68.4	65.3	71.4

(a) **Encoder type.** VideoMAC is compatible with both classical and modern ConvNets.

Dim.	$\mathcal{J} \& \mathcal{F}_m$	\mathcal{J}_m	\mathcal{F}_m
128	67.9	64.8	71.0
256	68.2	65.3	71.1
512	68.4	65.3	71.4
1024	68.1	65.0	71.2

(c) **Decoder width.** A decoder width of 512 delivers the best performance.

Blocks	$\mathcal{J} \& \mathcal{F}_m$	\mathcal{J}_m	\mathcal{F}_m
1	68.4	65.3	71.4
2	68.2	65.2	71.2
4	68.3	65.4	71.2
8	68.1	65.2	71.0

(b) **Decoder depth.** One block achieves competitive performance and efficiency.

Masking	$\mathcal{J} \& \mathcal{F}_m$	\mathcal{J}_m	\mathcal{F}_m
Asym.	66.5	63.4	69.6
Sym.	68.4	65.3	71.4

(d) **Masking design.** Symmetric (sym.) masking outperforms asymmetric (asym.) one.

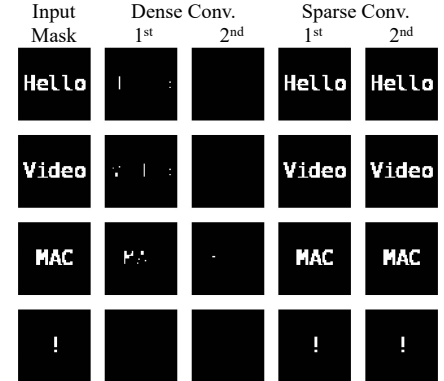


Figure 5. For masked modeling, dense convolution usually results in the dissipation of mask structures. The deployment of sparse convolution proves to be an effective solution, enabling ConvNet-based MIM / MVM.

Table 3. **VideoMAC encoder and decoder ablation experiments** on DAVIS17 [46]. We report ablation results for two of the most emblematic families of ConvNets. Unless specified otherwise, the encoder is CNXv2-S in default. ‘Dim’ indicates the number of decoder dimensions. Default settings employed in this paper are marked in blue.

Body Part Propagation. We evaluate VideoMAC on the Video Instance Parsing (VIP) dataset [73], and the quantitative results, measured by mIoU, are detailed in Table 1. This validation set comprises of 50 videos that concentrate on the propagation of 19 body parts, such as arms and legs, thereby necessitating a heightened level of precision in matching as compared to video object segmentation. We adhere to the same parameters as [25], adjusting the video frames to 560×560 dimensions. In the case of body part propagation, two variants of our VideoMAC demonstrate improvements by **3.1%** (CNXv2-S) and **6.3%** (RN50), respectively, in comparison with the top-performing SiamMAE [20] in MVM. Significantly, our RN50-based model outperforms CNXv2-S on the VIP dataset [73] due to the higher dimensionality of RN50 features (2048 vs. 768), which benefits complex tasks and improves accuracy in mask propagation.

Human Pose Tracking. We conduct the human pose tracking task on the validation set of JHMDB [26], which includes 268 videos and entails detecting 15 human keypoints. Conforming to the protocol delineated in [25], we resize video frames to a resolution of 320×320 , and employ

the Possibility of Correct Keypoints (PCK) metric [51] for quantitative evaluation. As shown in Table 1, both variants (*i.e.*, RN50 and CNXv2-S) of our VideoMAC outperforms all MVM methods in terms of PCK@0.1 and PCK@0.2.

Qualitative results. As depicted in Fig. 4, we present the qualitative results of mask propagation achieved by VideoMAC on three downstream tasks. Our VideoMAC yields notable performance across a spectrum of complex videos.

4.2. Ablation Studies

In this section, we conduct extensive ablation studies to analyze the component design of our VideoMAC.

Encoder design: hierarchical vs. isotropic. In this ablation study, we scrutinize diverse encoder architectures to our VideoMAC, *i.e.*, hierarchical and isotropic designs. Precisely, we adopted the ViT [14] design principles and restructured the CNXv2-S backbone from having a hierarchical structure into an isotropic one. This implies the creation of an architecture devoid of downsampling layers, ensuring consistent feature resolution (*e.g.*, 14×14) throughout all depths. As shown in Table 2, the performance dif-

Dataset	Size	$\mathcal{J}\&\mathcal{F}_m$	\mathcal{J}_m	\mathcal{F}_m	Augmentation	$\mathcal{J}\&\mathcal{F}_m$	\mathcal{J}_m	\mathcal{F}_m	Gap	$\mathcal{J}\&\mathcal{F}_m$	\mathcal{J}_m	\mathcal{F}_m
YT18 [68]	5.58 hours	68.4	65.3	71.4	spatial	68.4	65.3	71.4	1	68.4	65.3	71.4
K400 ¹ [27]	800 hours	67.8	65.2	70.3	color	64.9	61.8	68.0	5	66.4	63.5	69.2
					spatial + color	66.1	62.9	69.3	10	63.2	60.4	65.9

(a) **Training Dataset.** Our VideoMAC pre-trained on YT18 yields better performance and efficiency compared to K400. (b) **Data augmentation.** Minimal data augmentation achieves the best performance. (c) **Frame gap.** Video pairs with one frame gap perform the best results.

\mathcal{L}_c	\mathcal{L}_t	$\mathcal{J}\&\mathcal{F}_m$	\mathcal{J}_m	\mathcal{F}_m	Consistency	$\mathcal{J}\&\mathcal{F}_m$	\mathcal{J}_m	\mathcal{F}_m	Masking ratio	$\mathcal{J}\&\mathcal{F}_m$	\mathcal{J}_m	\mathcal{F}_m
\times	\times	51.1	52.2	50.0	$\gamma=0.0$	60.6	56.3	64.9	0.65	66.5	63.8	69.1
\times	\checkmark	60.6	56.3	64.9	$\gamma=0.1$	64.8	62.0	67.6	0.75	68.4	65.3	71.4
\checkmark	\times	63.4	63.4	68.7	$\gamma=0.5$	65.8	63.2	68.5	0.85	68.3	65.3	71.3
\checkmark	\checkmark	68.4	65.3	71.4	$\gamma=1.0$	68.4	65.3	71.4	0.95	60.3	59.7	60.9

(d) **Model loss.** VideoMAC utilizes consistency loss \mathcal{L}_c and target loss \mathcal{L}_t for the best. (e) **Reconstruction consistency.** When γ is set to 1.0, the consistency loss works the best. (f) **Masking ratio.** Our VideoMAC uses a masking ratio of 0.75 for the highest performance.

Table 4. **VideoMAC ablations** on data, loss, and masking ratio configuration. We report two official metrics \mathcal{J}_m and \mathcal{F}_m on DAVIS17 [46].



Figure 6. **Analysis of individual loss components** after VideoMAC pre-training 100 epochs on YouTube-VOS 2018 [68].

ference between our VideoMAC (*iso.*) and the ViT-based MVM methods (*e.g.*, VideoMAE [56] and DropMAE [65]) is marginal. By maintaining the hierarchical structure of ConvNets within VideoMAC models (*hie.*), we improve the performance significantly, even though their feature resolution (*i.e.*, 7×7) is merely half that of the ViT-based models [16, 20, 56, 65]. This can be ascribed to ConvNets’ capacity to preserve multi-scale spatial information.

Encoder backbone: classical and modern. For our study, we incorporate two emblematic ConvNet families, *i.e.*, ResNets [21] and ConvNeXts [36, 64]. As shown in Table 3a, we substitute dense convolution with sparse convolution in the baseline ConvNets. This adaptation seamlessly dovetails into our VideoMAC framework, thereby facilitating a ConvNet-based MVM. In addition, Fig. 5 highlights the influence of employing dense and sparse convolution on masks. It is apparent that sparse convolution effectively safeguards the structural integrity of the mask, thus enabling the integration of MAE pre-training within ConvNets.

Decoder design: blocks and dimensions. Adhering to a simplistic design principle, our VideoMAC adopts the decoder from the purely convolutional network structure deployed in ConvNeXt [36]. We conduct ablation experiments focusing on the depth and width of the decoder, with the results detailed in Tables 3b and 3c. The decoder exhibits competitive results with a single block and 512 dimensions, striking a trade-off between performance and efficiency.

Masking design: asymmetric vs. symmetric. In Table 3d, we investigate the impact of two masking strategies, namely




Figure 7. **Consistency loss** in relation to γ (left: identical encoder, CNXv2-S) and encoders (right: identical weight factor, $\gamma = 1.0$).

asymmetric (same masking ratio but different locations) and symmetric (same masking ratio and locations), on performance. Compared to asymmetric masking, symmetric masking simplifies the calculation of the reconstruction loss, as it covers the entire masked area. Moreover, it facilitates the inclusion of more reconstruction regions in the loss calculation, resulting in faster model convergence.

Data setting: dataset, augmentation, and gap. To assess the influence of different video datasets, Table 4a presents the results of pre-training our VideoMAC on both YT18 [68] and K400¹ [27]. The ablation results of VideoMAC on the two datasets differ by just **0.6%**, demonstrating the robustness and adaptability of our method to different datasets. For data augmentation, as shown in Table 4b, our VideoMAC achieves the best results using only spatial strategy, *i.e.*, random resized cropping and horizontal flipping. Additionally, we conduct an ablation study on the frame gap, and the optimal results are obtained with consecutive frames (*i.e.*, the gap is 1), as indicated in Table 4c.

The contribution of \mathcal{L}_c and \mathcal{L}_t . Previous MVM methods predominantly rely on a single loss, *i.e.*, the reconstruction loss, necessitating multiple frames to be simultaneously processed by the model. In contrast, our study utilizes a more computation-friendly strategy by inputting frame pairs into two separate encoders, followed by parameter updates using EMA. Consequently, we introduce two distinct loss terms: online \mathcal{L}_o and target \mathcal{L}_t reconstruction losses. Furthermore, to integrate temporal information, we

¹K400 is sampled to match YT18’s training data volume.




Figure 8. **Visualization of reconstruction examples** using a pre-trained VideoMAC (CNXv2-S) model with 0.75 masking ratio. These frames are randomly selected from the validation set of YT18 [68] and masked 60% of patches.

compute the loss between the reconstructed results of frame pairs, colloquially known as the reconstruction consistency loss \mathcal{L}_c . As shown in Table 4d, \mathcal{L}_t and \mathcal{L}_c improve **9.5%** and **12.3%** in terms of \mathcal{J} & \mathcal{F}_m , respectively, with the full VideoMAC model yielding the best results. The diverse loss components during pre-training are illustrated in Fig. 6.

The effect of weight factor γ . To investigate the impact of the weight factor on reconstruction consistency loss \mathcal{L}_c , we conduct experiments with different factor settings, *i.e.*, $\gamma \in \{0.0, 0.1, 0.5, 1.0\}$. As illustrated in Table 4e, superior performance is attained when γ approximates to 1. When γ equals 0, which signifies the absence of the reconstruction consistency loss term, a detrimental effect on performance is noticeable. Simultaneously, we record \mathcal{L}_c for various weight factors during pre-training, as depicted in Fig. 7. It becomes evident that an elevated value of γ results in a lesser loss. Ultimately, both the R50 and CNXv2-S models converge to comparable final losses. It is noteworthy that \mathcal{L}_c presents an initial decrease, succeeded by an increase, and ultimately reaches convergence, which aligns with the description in Eq. (2), implying that the ultimate reconstruction results between two frames converge to ϵ .

The impact of masking ratio. An ablation analysis of the masking ratio is conducted, and the results for our VideoMAC models (CNXv2-S) are documented in Table 4f. Ablation results portray a performance pattern that first ascends and subsequently descends with increasing masking ratios. Noticeably, the performance of our approach experiences a drastic decline at a 0.95 masking ratio caused by overfitting to the self-supervised pretext task.

4.3. Additional Experiments

VideoMAC for Image Recognition. As shown in Table 5, utilizing CNXv1-T [36] / CNXv2-T [64] pre-trained by our VideoMAC on YT18 [68], we achieve **82.3%** / **82.9%** accuracy on IN1K [11] image classification. In comparison to CNXv1-T and CNXv2-T trained from scratch, our approach improves the image recognition results by **0.2%** and **0.4%**. Remarkably, the performance garnered by VideoMAC (MVM) is proximate to or commensurate with the state-of-the-art ConvNet-based MIM methods [54, 64]. Given the pre-training cost (100 epochs vs. 1600 epochs) and the domain gap (video vs. image), we regard this im-

Method	Backbone	PT epoch	PT data	FT acc.
Supervised [36]	CNXv1-T	-	-	82.1
FCMAE [64]	CNXv1-T	1600	IN1K	82.3 (+0.2)
Spark [54]	CNXv1-T	1600	IN1K	82.4 (+0.3)
VideoMAC	CNXv1-T	100	YT18	82.3 (+0.2)
Supervised [64]	CNXv2-T	-	-	82.5
FCMAE [64]	CNXv2-T	1600	IN1K	83.0 (+0.5)
VideoMAC	CNXv2-T	100	YT18	82.9 (+0.4)

Table 5. **Comparing VideoMAC with previous ConvNet-based MIM approaches** (*e.g.*, Spark [54] and FCMAE [64]) on the image downstream task, *i.e.*, fine-tuning recognition accuracy (acc.) on IN1K [11]. ‘PT’ and ‘FT’ indicate pre-training and fine-tuning.

provement of our VideoMAC as noteworthy and promising compared to training from scratch and pre-training by MIM.

Visualization. We visualize reconstruction results from the YT18 [68] validation set to analyze performance in MVM pre-training. In Fig. 8, VideoMAC makes reasonable predictions about the masked regions, *e.g.*, color and contour.

5. Conclusion

In this study, we present VideoMAC, an MVM pre-training framework constructed fully using ConvNets. Our masked modeling for frame pairs is executed by introducing an on-line network (encoder and decoder) optimized by gradients, as well as the target network that updates parameters via EMA. Concurrently, for reconstruction results derived from both online and target networks, our proposed loss of reconstruction consistency between frame pairs paves the way for temporal modeling. Our VideoMAC empowers a multitude of ConvNets to reap the performance gains for downstream tasks (*e.g.*, video but also image) by MVM pre-training.

Limitation and Future Work. Existing ConvNet-based architectures typically use a fixed downsampling ratio, which in turn leads to a relatively fixed patch size. For future work, we will investigate ConvNets with adjustable patch size, allowing for higher feature resolution like ViT.

Acknowledgement. This work was supported by the National Natural Science Foundation of China (No. 62102182, 62202227, 62302217), Natural Science Foundation of Jiangsu Province (No. BK20220934, BK20220938, BK20220936), China Postdoctoral Science Foundation (No. 2022M721626, 2022M711635), Jiangsu Funding Program for Excellent Postdoctoral Talent (No. 2022ZB267), Fundamental Research Funds for the Central Universities (No.30923010303).

References

- [1] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Mike Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *ECCV*, pages 456–473, 2022. 1
- [2] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEiT: BERT pre-training of image transformers. In *ICLR*, 2022. 1, 2
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *NeurIPS*, pages 1877–1901, 2020. 1, 2
- [4] Sergi Caelles, Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. One-shot video object segmentation. In *CVPR*, pages 221–230, 2017. 2
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, pages 9912–9924, 2020. 2
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, 2021. 5
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020.
- [8] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *CVPR*, pages 15750–15758, 2021. 2
- [9] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *CVPR*, pages 3075–3084, 2019. 5
- [10] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. In *ICLR*, 2020. 1, 2
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 2, 5, 8
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019. 1, 2
- [13] Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. Unified language model pre-training for natural language understanding and generation. In *NeurIPS*, 2019. 2
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 1, 2, 6
- [15] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. In *CVPR*, pages 19358–19369, 2023. 1
- [16] Christoph Feichtenhofer, haoqi fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. In *NeurIPS*, pages 35946–35958, 2022. 1, 2, 3, 4, 5, 7
- [17] Peng Gao, Teli Ma, Hongsheng Li, Ziyi Lin, Jifeng Dai, and Yu Qiao. Mcmae: Masked convolution meets masked autoencoders. *NeurIPS*, 35:35632–35644, 2022. 2
- [18] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. In *NeurIPS*, pages 21271–21284, 2020. 2
- [19] Agrim Gupta, Stephen Tian, Yunzhi Zhang, Jiajun Wu, Roberto Martín-Martín, and Li Fei-Fei. Maskvit: Masked visual pre-training for video prediction. In *ICLR*, 2023. 2, 3
- [20] Agrim Gupta, Jiajun Wu, Jia Deng, and Li Fei-Fei. Siamese masked autoencoders. In *NeurIPS*, 2023. 1, 2, 3, 5, 6, 7
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 1, 2, 4, 5, 7
- [22] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 2, 3
- [23] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. 1, 2, 3, 5
- [24] Olivier J Hénaff, Skanda Koppula, Evan Shelhamer, Daniel Zoran, Andrew Jaegle, Andrew Zisserman, João Carreira, and Relja Arandjelović. Object discovery and representation networks. In *ECCV*, pages 123–143, 2022. 5
- [25] Allan Jabri, Andrew Owens, and Alexei Efros. Space-time correspondence as a contrastive random walk. In *NeurIPS*, pages 19545–19560, 2020. 2, 5, 6
- [26] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J Black. Towards understanding action recognition. In *ICCV*, pages 3192–3199, 2013. 1, 2, 5, 6
- [27] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 5, 7
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012. 2
- [29] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 2
- [30] Youngwan Lee, Jeffrey Willette, Jonghee Kim, Juho Lee, and Sung Ju Hwang. Exploring the role of mean teachers in self-supervised masked auto-encoders. In *ICLR*, 2023. 1, 5

- [31] Liulei Li, Tianfei Zhou, Wenguan Wang, Lu Yang, Jianwu Li, and Yi Yang. Locality-aware inter-and intra-video reconstruction for self-supervised correspondence learning. In *CVPR*, pages 8719–8730, 2022. 2
- [32] Liulei Li, Wenguan Wang, Tianfei Zhou, Jianwu Li, and Yi Yang. Unified mask embedding and correspondence learning for self-supervised video segmentation. In *CVPR*, pages 18706–18716, 2023. 2
- [33] Jihao Liu, Xin Huang, Jinliang Zheng, Yu Liu, and Hongsheng Li. Mixmae: Mixed and masked autoencoder for efficient pretraining of hierarchical vision transformers. In *CVPR*, pages 6252–6261, 2023. 4
- [34] Yuan Liu, Songyang Zhang, Jiacheng Chen, Zhaohui Yu, Kai Chen, and Dahua Lin. Improving pixel-based mim by reducing wasted modeling capability. In *ICCV*, pages 5361–5372, 2023. 1
- [35] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 2, 3
- [36] Zhuang Liu, Hanzhi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *CVPR*, pages 11976–11986, 2022. 1, 2, 3, 4, 7, 8
- [37] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *CVPR*, pages 3202–3211, 2022. 3
- [38] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017. 5
- [39] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 5
- [40] Xiankai Lu, Wenguan Wang, Martin Danelljan, Tianfei Zhou, Jianbing Shen, and Luc Van Gool. Video object segmentation with episodic graph memory networks. In *ECCV*, pages 661–679, 2020. 2
- [41] Xiankai Lu, Wenguan Wang, Jianbing Shen, Yu-Wing Tai, David J Crandall, and Steven CH Hoi. Learning video object segmentation from unlabeled videos. In *CVPR*, pages 8960–8970, 2020.
- [42] Gensheng Pei, Fumin Shen, Yazhou Yao, Guo-Sen Xie, Zhenmin Tang, and Jinhui Tang. Hierarchical feature alignment network for unsupervised video object segmentation. In *ECCV*, pages 596–613, 2022.
- [43] Gensheng Pei, Fumin Shen, Yazhou Yao, Tao Chen, Xian-Sheng Hua, and Heng-Tao Shen. Hierarchical graph pattern understanding for zero-shot video object segmentation. *TIP*, 32:5909–5920, 2023.
- [44] Gensheng Pei, Yazhou Yao, Fumin Shen, Dan Huang, Xingguo Huang, and Heng-Tao Shen. Hierarchical co-attention propagation network for zero-shot video object segmentation. *TIP*, pages 2348–2359, 2023. 2
- [45] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, pages 724–732, 2016. 5
- [46] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017. 1, 2, 5, 6, 7
- [47] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 2
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 2
- [49] Chaitanya Ryali, Yuan-Ting Hu, Daniel Bolya, Chen Wei, Haoqi Fan, Po-Yao Huang, Vaibhav Aggarwal, Arkabandhu Chowdhury, Omid Poursaeed, Judy Hoffman, Jitendra Malik, Yanghao Li, and Christoph Feichtenhofer. Hiera: A hierarchical vision transformer without the bells-and-whistles. *ICML*, 2023. 2
- [50] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 2
- [51] Jie Song, Limin Wang, Luc Van Gool, and Otmar Hilliges. Thin-slicing network: A deep structured model for pose estimation in videos. In *CVPR*, pages 4220–4229, 2017. 6
- [52] Thomas Tuytelaars, Tim Lebaillly, Behzad Bozorgtabar, Tinne Tuytelaars, and Jean-Philippe Thiran. Croc: Cross-view online clustering for dense visual representation learning. In *CVPR*, pages 7000–7009, 2023. 5
- [53] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, pages 6105–6114, 2019. 2
- [54] Keyu Tian, Yi Jiang, Qishuai Diao, Chen Lin, Liwei Wang, and Zehuan Yuan. Designing bert for convolutional networks: Sparse and hierarchical masked modeling. In *ICLR*, 2023. 2, 3, 4, 5, 8
- [55] Yunjie Tian, Lingxi Xie, Zhaozhi Wang, Longhui Wei, Xiaopeng Zhang, Jianbin Jiao, Yaowei Wang, Qi Tian, and Qixiang Ye. Integrally pre-trained transformer pyramid networks. In *CVPR*, pages 18610–18620, 2023. 4
- [56] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In *NeurIPS*, pages 10078–10093, 2022. 1, 2, 3, 4, 5, 7
- [57] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, pages 10347–10357, 2021. 2
- [58] Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. Videomae v2: Scaling video masked autoencoders with dual masking. In *CVPR*, pages 14549–14560, 2023. 2
- [59] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Lu Yuan, and Yu-Gang Jiang. Masked video distillation: Rethinking masked feature modeling for self-supervised video representation learning. In *CVPR*, pages 6312–6322, 2023. 2
- [60] Wenguan Wang, Xiankai Lu, Jianbing Shen, David J Crandall, and Ling Shao. Zero-shot video object segmentation via attentive graph neural networks. In *ICCV*, pages 9236–9245, 2019. 2
- [61] Wenguan Wang, Hongmei Song, Shuyang Zhao, Jianbing Shen, Sanyuan Zhao, Steven CH Hoi, and Haibin Ling.

- Learning unsupervised video object segmentation through visual attention. In *CVPR*, pages 3064–3074, 2019. [2](#)
- [62] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, pages 568–578, 2021. [2](#), [3](#)
- [63] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *CVPR*, pages 3024–3033, 2021. [5](#)
- [64] Sanghyun Woo, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. Convnext v2: Co-designing and scaling convnets with masked autoencoders. In *CVPR*, pages 16133–16142, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#), [8](#)
- [65] Qiangqiang Wu, Tianyu Yang, Ziquan Liu, Baoyuan Wu, Ying Shan, and Antoni B Chan. Dropmae: Masked autoencoders with spatial-attention dropout for tracking tasks. In *CVPR*, pages 14561–14571, 2023. [1](#), [2](#), [3](#), [5](#), [7](#)
- [66] Qiangqiang Wu, Tianyu Yang, Wei Wu, and Antoni B Chan. Scalable video object segmentation with simplified framework. In *ICCV*, pages 13879–13889, 2023. [2](#)
- [67] Z. Xie, Z. Zhang, Y. Cao, Y. Lin, J. Bao, and Z. Yao. Simmim: A simple framework for masked image modeling. In *CVPR*, pages 9653–9663, 2022. [1](#), [2](#)
- [68] Ning Xu, Linjie Yang, Yuchen Fan, Jianchao Yang, Dingcheng Yue, Yuchen Liang, Brian Price, Scott Cohen, and Thomas Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *ECCV*, pages 585–601, 2018. [5](#), [7](#), [8](#)
- [69] Yazhou Yao, Tao Chen, Guo-Sen Xie, Chuanyi Zhang, Fumin Shen, Qi Wu, Zhenmin Tang, and Jian Zhang. Non-salient region object mining for weakly supervised semantic segmentation. In *CVPR*, pages 2623–2632, 2021. [2](#)
- [70] Xiaosong Zhang, Yunjie Tian, Lingxi Xie, Wei Huang, Qi Dai, Qixiang Ye, and Qi Tian. Hivit: A simpler and more efficient design of hierarchical vision transformer. In *ICLR*, 2023. [2](#), [3](#)
- [71] Yurong Zhang, Liulei Li, Wenguan Wang, Rong Xie, Li Song, and Wenjun Zhang. Boosting video object segmentation via space-time correspondence learning. In *CVPR*, pages 2246–2256, 2023. [2](#)
- [72] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. Image bert pre-training with online tokenizer. In *ICLR*, 2022. [2](#)
- [73] Qixian Zhou, Xiaodan Liang, Ke Gong, and Liang Lin. Adaptive temporal encoding network for video instance-level human parsing. In *ACMMM*, pages 1527–1535, 2018. [1](#), [2](#), [5](#), [6](#)
- [74] Tianfei Zhou, Fatih Porikli, David J Crandall, Luc Van Gool, and Wenguan Wang. A survey on deep learning technique for video segmentation. *TPAMI*, 45(6):7099–7122, 2023. [2](#)
- [75] Xizhou Zhu, Yujie Wang, Jifeng Dai, Lu Yuan, and Yichen Wei. Flow-guided feature aggregation for video object detection. In *ICCV*, pages 408–417, 2017. [2](#)