

SIGMA: Sinkhorn-Guided Masked Video Modeling

Mohammadreza Salehi^{*}, Michael Dorckenwald^{*}, Fida Mohammad Thoker^{*‡},
Efstratios Gavves, Cees G. M. Snoek, Yuki M. Asano[†]

University of Amsterdam

Abstract. Video-based pretraining offers immense potential for learning strong visual representations on an unprecedented scale. Recently, masked video modeling methods have shown promising scalability, yet fall short in capturing higher-level semantics due to reconstructing predefined low-level targets such as pixels. To tackle this, we present Sinkhorn-guided Masked Video Modelling (SIGMA), a novel video pretraining method that jointly learns the video model in addition to a target feature space using a projection network. However, this simple modification means that the regular L2 reconstruction loss will lead to trivial solutions as both networks are jointly optimized. As a solution, we distribute features of space-time tubes evenly across a limited number of learnable clusters. By posing this as an optimal transport problem, we enforce high entropy in the generated features across the batch, infusing semantic and temporal meaning into the feature space. The resulting cluster assignments are used as targets for a symmetric prediction task where the video model predicts cluster assignment of the projection network and vice versa. Experimental results on ten datasets across three benchmarks validate the effectiveness of SIGMA in learning more performant, temporally-aware, and robust video representations improving upon state-of-the-art methods. Our project website with code is available at: <https://quva-lab.github.io/SIGMA>.

1 Introduction

Video-based pretraining offers an unprecedented scale of training data. With a simple estimate of 24 frames per second, 196 tokens per frame, and around 90K years of videos [53], YouTube alone would contain around 10,000 trillion tokens. Compare this to the language domain where, for example, large-scale pretraining such as Llama-2 [70] utilized “only” 2 trillion tokens. Videos offer not only a vast scale for pretraining, but also allow for understanding how objects interact with one another and transform across time. Such semantic temporal and spatial understanding is valuable for many applications ranging from autonomous driving [11, 36, 41] to robotic planning [27, 78], embodied AI [21], and learning world models [31, 46].

As in the language domain, the key to achieving scalability lies in utilizing self-supervised learning signals. Initial works on self-supervised video learning focused on solving pretext tasks such as temporal ordering [26] or playback speed prediction [85]. Those methods achieved relatively good performance on small-scale datasets like UCF-101, yet, these methods do not scale well with more pretraining data. More recently, the

^{*} equal contribution. [†] now at University of Technology Nuremberg. [‡] now at KAUST.

Vision Transformer (ViT) [19] has shown to scale well with data and parameter size [16, 61] and has been successfully adapted to videos [1, 23], albeit, initially requiring large-scale labeled data.

At present, self-supervised works utilizing a masked video modeling framework for vision transformers – operating similarly to the ‘filling-in-the-blank’ objective utilized in large language models – now dominate the field and obtain state-of-the-art performances [22, 25, 68]. In this framework, part of the input data is masked and the network is trained to reconstruct predefined targets such as pixel values [22, 68, 73], frame differences [45] or a Histograms of Oriented Gradients [25]. While scalable and simple, these current methods do not fully capture the higher-level temporal and spatial semantics, as we also show empirically in this paper. This is because individual patches or space-time tubes do not represent individual semantic units, while in the language domain, words or subwords do. Therefore, tasking the model to reconstruct them pushes it towards learning low-level features. Based on this, we propose a novel approach for masked video modeling that focuses on reconstructing more abstract temporal and spatial video representations.

In this work, we propose a new framework wherein the typically predefined reconstruction target space can be simultaneously learned alongside the video model. For this, a projection network is introduced which embeds both the visible and masked portions of the video, yielding deep feature reconstruction targets. However, employing a commonly used L2 reconstruction loss naively is ineffective due to the joint optimization of both networks leading to a trivial solution as both networks collapse to the same output irrespective of the input. To solve this, we introduce SIGMA: Sinkhorn-guided masked video modeling, where deep features of space-time tubes are regularised by optimal transport uniformly across clusters. This effectively acts as a high-entropy regularization constraint and enforces similar space-time tube features to be assigned to the same centroid, infusing semantic meaning into the feature space. These cluster assignments and centroids are learned in an online manner using the fast Sinkhorn-Knopp algorithm, yielding feature pseudo-labels as targets. With these targets, we formulate our loss objective as a symmetric prediction task, where the features from each branch – the video model and the projection network – cross-predict the cluster assignment of the other. By doing so, we force the features of space-time tubes to be expressed by a limited number of clusters, enforcing semantic-rich concepts, while eliminating the dependency on predefined targets such as the masked pixel values, commonly used in prior works. Moreover, despite our cross-prediction task, we do not rely on any augmentations [18, 52, 55] or crops [5, 63], making our model stable and easy to train.

We conduct experiments on ten datasets across three benchmarks and demonstrate the superiority of our approach compared to the baseline VideoMAE [68], as well as

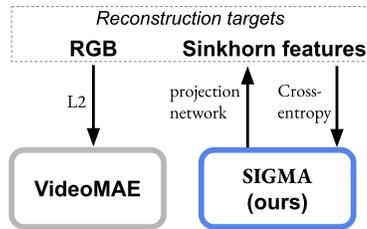


Fig. 1. Overview of our idea. Compared to VideoMAE, which uses RGB pixels as targets, we generate Sinkhorn-regularised features as reconstruction targets. This obtains more semantic features and yields better pre-training performance.

current state-of-the-art methods. We show that in a frozen evaluation setting our method outperforms previous works considerably on various video and even image datasets showcasing better transferability. Moreover, we evaluate the object-level understanding of our method, by evaluating its performance on an unsupervised video object segmentation benchmark [58], demonstrating better spatial and temporal semantics. Finally, we evaluate the generalization abilities on the SEVERE benchmark [65] and improve upon existing works. These findings confirm that SIGMA learns more performant, temporally-aware and robust representations. We will make the code and models available.

Our contributions can be summarized as follows:

- We propose SIGMA: a masked video modeling framework in which the reconstruction target space is jointly learned along with the video model.
- We extensively evaluate our pretraining method across standardized sets of datasets and benchmarks and obtain state-of-the-art results.
- We highlight the flexibility of our approach by including DINO-pretrained features as reconstruction targets, which obtains further gains.

2 Related Works

Self-supervised Video Representation Learning. Self-supervised learning aims to learn features that can discriminate within an input distribution without the need for supervision. This is often achieved in images, by training models to accomplish pretext tasks, as outlined in works like [49], [28], and [87], or instance-discrimination tasks, detailed in [3, 8, 13, 50, 79]. The underlying premise is that the features developed via this methodology are broadly applicable, as they can differentiate between a large number of instances or solve general pretext tasks. In videos, a similar approach is employed, with an increased focus on temporal aspects. Consequently, numerous pretext tasks have been proposed, including spatio-temporal puzzle solving [37], pace prediction [71], Odd-One-Out [26], temporal order verification [47, 81], speed prediction [6, 12], sorting sequences [42], playback rate perception [85], and ranking transformations [20]. For instance-discrimination, the methods such as [15, 55, 64, 67, 72, 84] have been proposed, which set out to learn a representation that can separate between different instances, while ignoring the meaningless variances introduced by spatio-temporal data augmentations [13, 18, 33, 50, 66]. Unlike these methods, our approach does not depend on spatio-temporal augmentations. This enhances its generalizability and scalability, as it eliminates the need for dataset-specific augmentation design.

Masked Input Modeling for Vision Transformers. The idea of reconstructing masked input as a generative pretext task was first introduced by [51]. MAE [32] showed that the features learned by this approach are very general when the backbone is replaced by vision transformers [19]. Since then, lots of methods have been proposed that expand on the idea, such as [25, 48, 74, 77, 80] in the image domain. Recent research, including JEPA [4], has proposed the novel approach of mask feature prediction within the latent space, demonstrating improved performance on various image benchmarks, at the expense of instability challenges. VideoMAE [69] employs the same approach for the first

time in video. Follow-up works have tried to improve VideoMAE by predicting low-level predefined targets instead of pixels [74, 75] or improve the temporal understanding of VideoMAE by either designing masks that mainly focus on the moving objects or explicitly adding motion information in the objective function [34, 63]. MME [63], for instance, applies two major modifications to VideoMAE. First, following [25], the decoder predicts HOG features of the masked area instead of the pixel values. Second, it extracts object motion trajectories using optical flow and predicts these trajectories in conjunction with the masked features using the decoder part. However, as motion trajectories become inaccurate in the presence of camera motion, the entire training dataset undergoes preprocessing to eliminate the effects of camera motion before training. By similar intention of learning better motion information, MGM [34] and MGMAE use the hand-crafted motion vectors used by H.264 codec [57] or optical flow modality to mask tubes representing moving objects more than other areas. Unlike such methods, our method is general and does not require any preprocessing or hand-crafted features. Instead, we propose to project spatio-temporal tubes in a deep feature space as the targets of the decoder to jointly learn high-level motion and appearance patterns.

3 Methodology

In Sec. 3.1, we first review the common approach to masked video modeling, which focuses on the precise reconstruction of masked pixel values, leading to the extraction of low-level features. As a solution, we introduce a projection network φ in Sec. 3.2 that is jointly optimized with the video model and whose features are used as targets for the reconstruction task. However, this adjustment leads to a trivial solution. To circumvent this, we employ a cluster bottleneck strategy, as detailed in Section 3.3. Features of similar space-time tubes are matched to the same cluster which infuses semantic meaning into our feature space. With these pseudo-labels, we can formulate a self-supervised prediction task with which we optimize the parameters of our networks. An overview of our method is depicted in Figure 2.

3.1 Masked video modeling

In masked video modeling [22, 45, 63, 68, 73], the primary objective is the precise reconstruction of masked pixel values. Specifically, a video V , composed of frames $\{x_1, \dots, x_T\}$, is segmented into a set of space-time tubes \mathcal{T} with a significant proportion (approximately 90%) being masked motivated by the inherent redundancy in video data. A video model Ψ with an asymmetric encoder-decoder architecture is trained to accurately reconstruct the pixels of the masked space-time tubes \mathcal{T}^M leveraging the unmasked tubes $\mathcal{T}^N = \mathcal{T} \setminus \mathcal{T}^M$ and the positional information of the masked patches via

$$\mathcal{L}_2 = \frac{1}{N} \sum_{i=1}^N \|\mathcal{T}_i - \Psi(\mathcal{T}^N)_i\|_2^2, \quad (1)$$

with N the number of masked space-time tubes, \mathcal{T}_i denotes the pixels of the i -th space-time tube, and $\Psi(\mathcal{T}^N)_i$ represents the reconstructed pixels for this tube given the set

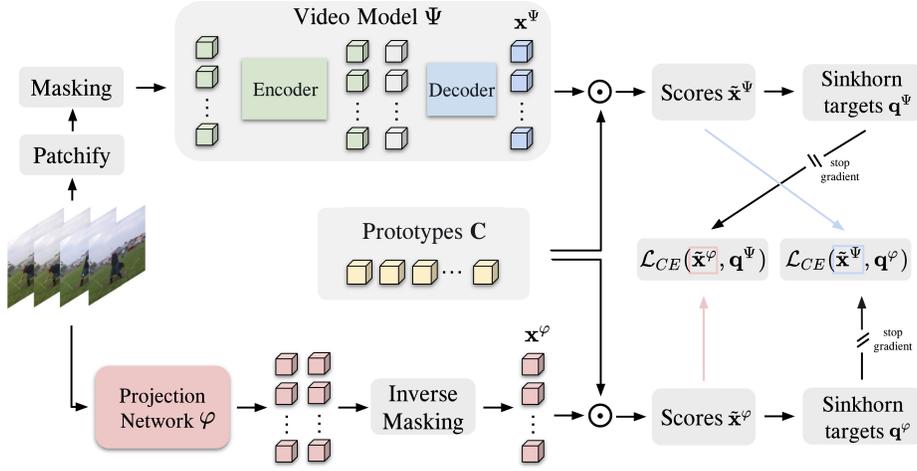


Fig. 2. Overview of our proposed method SIGMA . A given video is embedded with the projection network φ leading to features \mathbf{x}^φ . The video model Ψ predicts feature embeddings \mathbf{x}^Ψ of the masked space-time tubes. Both embeddings are projected onto the learnable prototypes representing cluster centroids. Cluster assignments are created with an adapted Sinkhorn algorithm enforcing equipartition across all prototypes. These pseudo-labels are then used as targets for the predictive task \mathcal{L}_{CE} with which the networks are optimized.

of all non-masked tubes \mathcal{T}^N . This loss function encourages strictly predicting the pixel values of the input video. Thus, mainly low-level features are captured in the trained video model Ψ as evidenced by low linear probing performance (see Tab. 1). Yet, video reasoning necessitates an understanding of how scenes and objects interact with one another and transform across time. To achieve such a higher-level semantic understanding, we first exchange the target space in which the reconstruction loss is applied, to a deeper feature space.

3.2 From pixel to feature reconstruction

Rather than directly predicting pixel values, we change the target space to a feature space by introducing a projection network that can be jointly optimized. This enables the projection network to autonomously learn the target space, diverging from previous works that use predefined targets such as frame differences [83], motion trajectories [63] or HOG features [25]. The projection network φ embeds space-time tubes from a given video into features representing $\mathbf{x}^\varphi = \varphi(\mathcal{T})$ (see bottom stream in Fig. 2). We only select the features from the space-time tubes \mathcal{T}^M that are used for the masked prediction task in the video model and name this sampling process inverse masking in Fig. 2. Thus, the target space has been replaced, and Eq. 1 changes to

$$\mathcal{L}_2^F = \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i^\varphi - \mathbf{x}_i^\Psi\|_2^2, \quad (2)$$

with the prediction $\mathbf{x}^\Psi = \Psi(\mathcal{T}^N)$. However, this leads to a trivial solution as both, the parameters of the projection network φ and the encoder-decoder video model Ψ , are

being jointly optimized, and can agree to map all space-time tubes to the same feature vector $\varphi(\mathcal{T}^N) = \Psi(\mathcal{T}^N)$.

3.3 SIGMA: Sinkhorn-Guided Masked Video Modeling

To circumvent this trivial solution, we constrain the feature space to be partitioned into a limited number of clusters. This can be interpreted as a bottleneck with enforced high entropy. We accomplish this in an online manner by mapping the features \mathbf{x} of space-time tubes to a set of learnable prototypes $\mathbf{C} = \{\mathbf{c}_1, \dots, \mathbf{c}_K\}$, representing cluster centroids. For that, we define a mapping $\mathbf{Q} \in \mathbb{R}_+^{K \times B}$ of embeddings $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_B]$ to prototypes \mathbf{C} . An equipartition constraint is introduced on \mathbf{Q} to enforce that all prototypes are equally used, mitigating the trivial solution of all data points collapsing onto a single prototype. Due to the limited number of prototypes, similar and nearby space-time tubes are assigned to the similar prototypes, infusing semantic spatial and temporal meaning into the feature space. This assignment can be formulated as an entropy-regularised optimal transport problem [2] for all space-time tubes over a mini-batch \mathcal{B}

$$\min_{\mathbf{Q}} \langle \mathbf{Q}, -\log \mathbf{X} \rangle + \frac{1}{\lambda} \text{KL}(\mathbf{Q} \| r c^\top), \quad (3)$$

$$\text{with } r = \frac{1}{K} \cdot \mathbb{1}, \quad c = \frac{1}{B} \cdot \mathbb{1}. \quad (4)$$

Here λ controls the entropy regularisation, r, c are the marginals for the prototypes, and $B = N_B \times N_T$ being all samples N_B in the mini-batch \mathcal{B} for all space-time tubes N_T . We rely on the Sinkhorn algorithm [14] to solve Eq. 4, which can be done extremely quickly on the GPU [2, 9] and yields soft pseudo-labels \mathbf{q} , such that $\text{argmax}(\mathbf{q}) = \mathbf{Q}$. With that, pseudo-labels \mathbf{q}^φ are generated for the embeddings \mathbf{x}^φ of the projection network. These pseudo-labels are then used as targets for the video model Ψ and vice versa. Then the goal of φ and Ψ is to predict each other cluster assignments \mathbf{q} based on the projection of the features onto the shared prototypes [9]. We formulate this predictive task between the pseudo-label \mathbf{q} and the scores $\tilde{\mathbf{x}} = \mathbf{x}^\top \mathbf{c}$ using cross-entropy

$$\mathcal{L}_{CE}(\tilde{\mathbf{x}}, \mathbf{q}) = - \sum_k \mathbf{q}_{(k)} \log \mathbf{p}_{(k)} \quad (5)$$

$$\mathbf{p}_{(k)} = \frac{\exp\left(\frac{1}{\tau} \tilde{\mathbf{x}}_k\right)}{\sum_{k'} \exp\left(\frac{1}{\tau} \tilde{\mathbf{x}}_{k'}\right)} \quad (6)$$

where τ represents the temperature parameter. We use softmax to define a probability distribution $\mathbf{p}_{(k)}$ between \mathbf{x} and all K prototypes in \mathbf{C} . The loss function for all space-time tubes \mathcal{T} in a mini-batch \mathcal{B} can be written as a symmetric prediction task

$$\mathcal{L} = \frac{1}{B} \sum_{i=1}^B [\mathcal{L}_{CE}(\tilde{\mathbf{x}}_i^\varphi, \mathbf{q}_i^\Psi) + \mathcal{L}_{CE}(\tilde{\mathbf{x}}_i^\Psi, \mathbf{q}_i^\varphi)], \quad (7)$$

with $\mathcal{L}_{CE}(\tilde{\mathbf{x}}^\varphi, \mathbf{q}^\Psi)$ the loss for predicting the pseudo-labels based on the feature of the video model Ψ and $\mathcal{L}_{CE}(\tilde{\mathbf{x}}^\Psi, \mathbf{q}^\varphi)$ the one from the projection network φ . By using feature pseudo-labels as targets for our symmetric loss, we solve the shortcomings

of low-level predefined targets in Eq. 1 and circumvent the trivial solution for deep targets outlined in Eq. 2. In contrast to prior works [5, 18, 52, 55, 56, 67] that rely on handcrafted augmentations to generate different views for their loss objective, our loss Eq. 7 circumvents that by instead predicting deep feature targets of masked space-time tubes. The joint space defined by the Sinkhorn clustering also mitigates the need for a momentum encoder as required by previous works [5, 18, 55] and thus allows us to vary the architecture choice for our projection network φ , as we will demonstrate in our experiments.

4 Experiments

We evaluate our method on a total of ten different datasets and across three common benchmark settings. In Sec. 4.1 we first compare our approach against state-of-the-art video models in a linear probing (frozen backbone) and the standard full finetuning setting. Then, in Sec. 4.2 we benchmark the semantic spatial and temporal understanding by reporting unsupervised semantic segmentation performance and visualizing the segmentation masks. In Sec. 4.3 we evaluate our approach on the SEVERE benchmark [65, 67] specifically designed to analyze the generalization performance of video models. Lastly, we ablate parts of our approach to give more insights in Sec. 4.4.

Implementation details. As a projection network φ we use a simple 3-layer Multi-Layer Perceptron (MLP) consisting of a linear layer with 1024 neurons followed by a GELU activation function and a final linear layer mapping to the feature dimension of 256 for the predictive task. This MLP embeds a given space-time tube and is trained jointly with our video model ϕ . This variant is abbreviated as SIGMA w/ MLP. Moreover, our approach allows us to leverage pretrained models such as DINO [10] which is kept frozen while training the video model (SIGMA w/ DINO). We pretrain these variants on: **Something-Something V2 (SSv2)** [30] and **Kinetics-400 (K400)**. SSV2 contains 220K videos with 174 action classes and is considered motion-heavy because of its focus on motion and directional aspects inherent to the actions. K400 is acknowledged as the benchmark for video recognition evaluations, encompassing 240K Internet-sourced videos classified into 400 action categories. We follow previous works for pretraining and use a temporal stride τ of 2 for SSV2 and 4 for K400. We sample a clip of 16 frames from the raw video at a resolution of 224×224 pixels. We use space-time tube embeddings like in VideoMAE [68] where each cube of size $2 \times 16 \times 16$ is treated as one token embedding and are extracted with a 3D convolution layer. We also use the commonly high masking ratio of 90% with which tubes are randomly masked. To be comparable with previous works and use the same configurations for the small and base ViT backbones. We follow [68] to use AdamW [44] optimizer with a base learning rate $1.5e^{-4}$, weight decay 0.05, $\beta = [0.9, 0.95]$, and cosine learning rate decay. For downstream tasks, we employ average pooling upon all output tokens before a final layer is added to map from the feature dimension to the number of classes used for classification. We follow the full finetuning setup from VideoMAE [68] and list all hyperparameters for pretraining and finetuning in the supplemental material. The source code of this project will be released upon acceptance.

Table 1. Benchmark I: Frozen evaluation of masked video modeling methods. A linear layer on top of the frozen ViT-B backbone is optimized. The ViT-B backbones are pretrained on Kinetics-400 (K400). We evaluated the official publicly released model from the corresponding method. SIGMA consistently outperforms previous masked video modeling works across all video and image datasets, often considerably.

	SSv2	K400	UCF	HMDB	IN-1K	C-100
MVD [76]	12.2	18.7	49.1	28.6	20.0	34.1
MME [63]	16.6	19.1	56.0	37.1	15.8	34.2
VideoMAE [68]	17.5	20.7	58.6	37.7	20.2	40.4
MGM [35]	16.8	24.9	64.4	41.3	20.4	44.1
SIGMA w/ MLP	<u>19.9</u>	<u>30.7</u>	<u>73.8</u>	<u>45.0</u>	<u>24.1</u>	<u>46.2</u>
SIGMA w/ DINO	20.8	47.5	80.7	52.3	45.0	66.7

4.1 Benchmark I: Comparisons for linear and full finetuning

We conduct the linear and full finetuning experiments on four commonly used datasets including **SSv2**, **K400**, **UCF-101** [62] and **HMDB-51** [39]. We refer to the supplementary material for details on the datasets. In addition, we also show the effectiveness of video models on image datasets such as **ImageNet-1K** (IN-1K) [17], and **CIFAR-100** (C-100) [38].

Frozen evaluation setup. We analyze the downstream abilities of our approach and compare it against state-of-the-art methods using a frozen evaluation setting. The underlying ViT-B backbone pretrained on K400 remains frozen while a linear layer on top is optimized for the specific target dataset. This frozen setup better evaluates the pretrained models compared to full finetuning where all the learned parameters are overwritten. We compare against the recent masked video modeling works MGM [22], MME [63], MVD [76], and VideoMAE [68]. We used the official publicly released pretrained models from those methods and used the commonly used 800 epochs checkpoint except for MVD which pretrain for 1600 epochs and distill for 400 epochs.

Frozen evaluation results. From the results in Tab. 1 we observe a big improvement in performance with our method against all state-of-the-art methods while using only a simple MLP as a projection network leading to an average improvement of 4.5% across all datasets and domains. Furthermore, an even bigger improvement is observed when we train our method using DINO as a projection network with an average improvement of 16.8% across all datasets. Surprisingly, we observe that MVD, which relies on predefined MAE feature targets, achieves the lowest performance, showcasing that the ability of our method to jointly learn a more semantic target space is crucial. Our method outperforms VideoMAE and MGM, which use pixel values as reconstruction targets, highlighting the efficacy of our Sinkhorn-guided approach that predicts feature cluster assignments. In summary, these findings confirm that our approach leads to more spatial and temporal semantic features.

Full finetuning setup. In this experiment, we evaluate the downstream action classification abilities of our method in a full finetuning setup for SSv2 and K400. We use the ViT-S and ViT-B backbones and pretrain for 800 epochs on the respective datasets following previous works [22, 45, 68]. For clarity, we compare against methods with

Table 2. Benchmark I: Comparison for full finetuning on Something-Something V2 (SSv2). The top part comprises supervised methods while the remaining methods are pretrained in a self-supervised manner. The middle section evaluates models trained on Kinetics 400 (K400) data for pretraining whereas the bottom part mainly used SSv2 data. We compare against previous methods pretrained on the ViT-Base backbone for 800 epochs. A full table with all previous works using also different pretraining setups is provided in the supplemental. M. Guid. denotes motion guidance such as optical flow used e.g. reconstructing targets or masking. Our method achieves state-of-the-art performance on SSv2 w/ and w/o motion guidance.

	Method	M. Guid.	Backbone	Epochs	Extra data	Frames	Params	Top-1
<i>supervised baselines</i>								
	SlowFast [24]	–	ResNet101	–	K400	8+32	53	63.1
	MViTv1 [23]	–	MViTv1-B	–	K400	64	37	67.7
	TimeSformer [7]	–	ViT-B	–	IN-21K	8	121	59.5
	VideoSwin [43]	–	Swin-B	–	IN-21K	32	88	69.6
<i>self-supervised</i>								
K400 pretraining	SIGMA-DINO (ours)	×	ViT-S	800	IN-1K+K400	16	22	68.7
	OmniMAE [29]	×	ViT-B	800	IN-1K+K400	16	86	69.0
	VideoMAE [68]	×	ViT-B	800	K400	16	87	68.5
	MME [63]	✓	ViT-B	800	K400	16	87	<u>70.5</u>
	SIGMA-MLP (ours)	×	ViT-B	800	K400	16	87	69.8
	SIGMA-DINO (ours)	×	ViT-B	800	IN-1K+K400	16	87	71.1
SSv2 pretraining	VideoMAE [68]	×	ViT-S	2400	–	16	22	66.8
	SIGMA-DINO (ours)	×	ViT-S	2400	IN-1K	16	22	68.6
	OmniMAE [29]	×	ViT-B	800	IN-1K	16	86	69.5
	VideoMAE [68]	×	ViT-B	800	–	16	87	69.6
	CMAE-V [45]	×	ViT-B	800	–	16	87	69.7
	SIGMA-MLP (ours)	×	ViT-B	800	–	16	87	70.4
	SIGMA-DINO (ours)	×	ViT-B	800	IN-1K	16	87	70.9
	MME [63]	✓	ViT-B	800	–	16	87	70.0
	MGM [22]	✓	ViT-B	800	–	16	87	70.6
	MGMMAE [35]	✓	ViT-B	800	–	16	87	71.0
	SIGMA-MLP (ours)	✓	ViT-B	800	–	16	87	71.2
	SIGMA-DINO (ours)	✓	ViT-B	800	IN-1K	16	87	71.2

this commonly used pretraining setup and report a table featuring all previous works in the supplemental material.

Full finetuning results. We report results for SSv2 in Tab. 2 and for K400 in Tab. 3. Given the nature of full finetuning in which all learned parameters of the network are optimized the gap between methods is smaller. Still, we observe a substantial improvement of on average $\sim 1\%$ across all datasets and backbones from our method using a simple MLP as a projection network compared to VideoMAE [68]. This confirms the importance of utilizing a learnable target reconstruction space rather than simply regressing pixel values. SIGMA can also incorporate motion guidance via advanced masking following MGMMAE [35], further improving the results on SSv2, especially for the MLP variant. Our method with DINO as a projection network achieves state-

Table 3. Benchmark I: Comparison for full finetuning on Kinetics 400 (K400). We compare against previous methods for pretraining the ViT-Base backbone for 800 epochs on K400 and subsequently, fully finetuning the backbone with the K400 labels. A full table with all previous methods using also different setups is provided in the supplemental. M. Guid. denotes motion guidance such as optical flow used e.g. reconstructing targets or masking. Our method achieves state-of-the-art performance on K400.

	Method	M. Guid.	Backbone	Epochs	Extra data	Frames	Params	Top-1
	<i>supervised</i>							
	SlowFast [24]	×	ResNet101	–	–	16+64	60	79.8
	MViTv1 [23]	×	MViTv1-B	–	–	32	37	80.2
	TimeSformer [7]	×	ViT-B	–	IN-21K	96	430	80.7
	VideoSwin [43]	×	Swin-L	–	IN-21K	32	197	83.1
	<i>self-supervised</i>							
	VideoMAE [68]	×	ViT-S	1600	–	16	87	79.0
	SIGMA-DINO (ours)	×	ViT-S	800	IN-1K	16	87	79.4
K400 pretraining	VideoMAE [68]	×	ViT-B	800	–	16	87	80.0
	OmniMAE [29]	×	ViT-B	800	IN-1K	16	87	80.8
	CMAE-V [45]	×	ViT-B	800	–	16	87	80.2
	MGM [22]	✓	ViT-B	800	–	16	87	80.8
	MGMMAE [35]	✓	ViT-B	800	–	16	87	<u>81.2</u>
	SIGMA-MLP (ours)	×	ViT-B	800	–	16	87	80.2
	SIGMA-DINO (ours)	×	ViT-B	800	IN-1K	16	87	81.6

of-the-art performance on SSv2 and K400. We even outperform the VideoMAE ViT-B model with our ViT-S K400 pretrained model on SSv2.

4.2 Benchmark II: Unsupervised video object segmentation

Setup. To assess the temporal and spatial semantics learned by our proposed method, we employ the unsupervised video object segmentation benchmark introduced by [58]. This benchmark evaluates the pretrained video encoder’s ability to generate temporally consistent object segmentation maps through its space-time features. In this regard, it deviates from the other reported evaluation settings where dense space-time features are pooled for a global clip representation. Instead, space-time features are grouped into clusters via k-means, based on a predefined number of clusters denoted by the K parameter. These clusters are subsequently matched against ground truth object segmentation maps using the Hungarian algorithm [40], and their overlap (mIoU) is calculated to determine the unsupervised object segmentation score. This procedure is called clustering when K matches the count of objects in the ground truth video, and over-clustering when K exceeds this count. For both procedures, we report the segmentation performance in terms of mIoU for two datasets: **DAVIS** [54] and **YTVOS** [82]. Please see the supplementary materials for details on the datasets.

Results. As shown in Tab. 4, our method with a simple MLP as a projection layer outperforms almost all the competitors by 1.5% on average. SIGMA with DINO as a projection network works even better and outperforms all previous works by a large

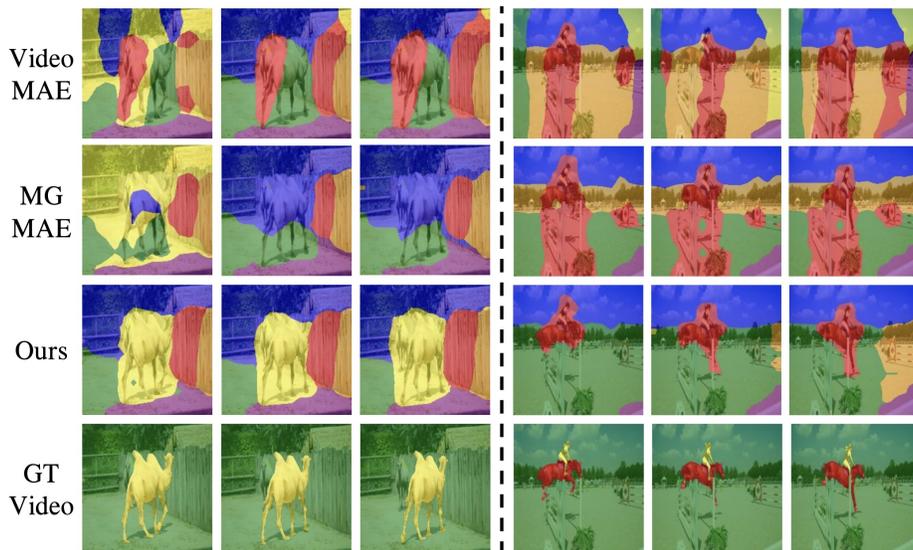


Fig. 3. Benchmark II: Unsupervised video object segmentation results on DAVIS. We visualize the abilities of masked video modeling methods to produce temporally consistent semantic segmentation masks. SIGMA provides more coherent and consistent object cluster maps compared to other methods. This shows that our learned features have better temporal and spatial understanding.

margin across all metrics by approximately 3% on average using the ViT-B backbone and by approximately 4.5% with ViT-S. These findings indicate that our model produces more semantic space-time representations. Moreover, although MGM [34] concentrates on developing masks for moving objects to enhance feature representation, our method still achieves a 2% to 5% improvement without changing the commonly used random tube masking. This result emphasizes the efficacy of SIGMA over methods based on pixel-level reconstruction for dense downstream tasks. In Fig. 3 we qualitatively compare the cluster maps produced using our encoder’s features compared to VideoMAE [69] and MGM [34] on DAVIS. For that, each clip is passed to the encoder and its tube representations are extracted. Then, the extracted features are clustered by k-means with a K higher than object counts specified by the ground truth. The reason is, for DAVIS the number of ground truth objects is usually smaller than the real number of objects. Hence, overclustering is a more realistic setting. Finally, the extracted cluster maps are resized to match the input size and overlaid on the input. SIGMA with DINO as a projection network produces accurate segmentations for each video clip, indicating a better understanding of the object-level concepts. This can be seen for example in the left visualization in Fig. 3 where the camel is well segmented, compared to VideoMAE and MGM. Additionally, our learned features exhibit better consistency over time than the temporal flickering observed in VideoMAE and MGM. These results indicate that our method leads to more robust video representations with object temporal consistency.

Table 4. Benchmark II: Unsupervised video object segmentation. We follow the evaluation protocol from [58] and report mIoU for clustering and overclustering. SIGMA consistently achieves better results compared to other methods across different backbones and datasets. This shows that our method learns better semantic and temporal consistent features.

Method	Backbone	Clustering		Overclustering	
		YTVOS	DAVIS	YTVOS	DAVIS
VideoMAE [68]	ViT-S	33.7	28.3	53.9	52.8
MVD [76]	ViT-S	32.9	27.0	50.3	43.3
SIGMA-DINO (ours)	ViT-S	38.2	30.6	60.0	56.8
VideoMAE [68]	ViT-B	34.1	27.2	49.7	50.9
MVD [76]	ViT-B	33.1	29.5	52.9	45.2
MGMAE [35]	ViT-B	34.5	28.2	55.4	54.8
SIGMA-MLP (ours)	ViT-B	<u>35.9</u>	<u>30.2</u>	<u>57.4</u>	52.9
SIGMA-DINO (ours)	ViT-B	36.4	30.6	60.4	56.5

4.3 Benchmark III: SEVERE generalization

Setup. To investigate the generalization capability of self-supervised video representations learned by our model we evaluate it on the challenging SEVERE benchmark introduced in [65] and extended by [66]. This benchmark comprises eight experiments focused on four downstream generalization factors: *domain shift*, *sample efficiency*, *action granularity*, and *task shift*. Domain shift is assessed using **Something-Something v2** and **FineGym(Gym99)** [59], both differing in domain relative to pre-training dataset Kinetics-400. Sample efficiency is tested through low-shot action recognition on **UCF101** and **FineGym**, with only 1,000 training samples available for fine-tuning. Action granularity is explored by evaluating semantically similar actions using subsets **FX-S1** and **UB-S1** from FineGym, where action classes pertain to the same element of a gymnastic routine (e.g., FX-S1 representing types of jumps). Task shift assesses performance beyond single-label action recognition, involving temporal repetition counting on UCFRep [86], a subset of UCF101, and multi-label action recognition on **Charades** [60]. Detailed experimental setups for each subset are provided in the supplementary.

Results. We follow the same setup as in the original SEVERE benchmark [65, 66]. In particular, we compare with VideoMAE [68] and SVT [56] as reported in [66] and extend the benchmark with SIGMA, MVD [76], and MGM [22]. All models are pre-trained on Kinetics-400 using the ViT-B architecture. Results are shown in Table 5. **Domain Shift.** Our approach SIGMA (both with MLP and DINO) outperforms all prior methods on Gym99, e.g. +7.2% over MVD and +3.1% over VideoMAE, demonstrating that the representation learned by our model is robust to domain shift in the downstream dataset. This is also demonstrated by strong SSv2 fine-tuning performance with K400 pretraining, as shown in Table 2. **Sample Efficiency.** For sample efficiency, we achieve a considerable gain over all prior works on Gym (10^3), e.g., +2.1% over VideoMAE [68] and +4.0% over MGMAE [34] showing the effective adaptability of our model to low shot action recognition. **Action granularity.** For fine-grained actions in

Table 5. Benchmark III: SEVERE Generalization [65]. We evaluate masked video modeling methods for generalizability in domain shift, sample efficiency, action granularity, and task shift following [65]. SIGMA achieves strong generalization performance outperforming prior works across all configurations. We use the original severe codebase [65, 67] to evaluate the publicly available models for all the methods.

	Domains		Samples (10^3)		Actions		Tasks		Mean
	SSv2	Gym99	UCF	Gym	FX-S1	UB-S1	UCF-RC↓	Charades	
SVT [56]	59.5	62.3	<u>83.9</u>	18.5	35.4	55.1	0.421	35.5	51.0
MVD [76]	<u>70.0</u>	82.5	66.7	17.5	31.3	50.5	0.184	16.1	52.1
VideoMAE [69]	68.6	86.6	74.6	25.9	42.8	65.3	<u>0.172</u>	12.6	57.4
MGMAE [35]	68.9	87.2	77.2	24.0	33.7	79.5	0.181	17.9	58.8
SIGMA-MLP (ours)	<u>69.8</u>	<u>87.4</u>	80.2	<u>26.8</u>	<u>46.0</u>	<u>79.7</u>	0.178	20.1	<u>61.5</u>
SIGMA-DINO (ours)	70.9	89.7	84.1	28.0	55.1	79.9	0.169	<u>23.3</u>	64.3

FX-S1 and UB-S1, our method achieves the best performance as well, with a considerable improvement over other models, *e.g.*, +12.3% and +14.6% over VideoMAE. These results demonstrate that the video representation learned by our method is better suited to fine-grained actions than existing self-supervised methods. **Task Shift.** For the task shift to repetition counting our method achieves the best performance. For multi-label action recognition on Charades, our approach lags behind SVT which employs a contrastive learning paradigm, however, we outperform all masked video modeling methods. **Overall SEVERE Performance.** Finally, we compare the mean across all generalizability factors. Our method has the best mean performance when using DINO as the projection network and achieves the second-best performance when employing MLP as the projection network. We conclude our method improves the generalizability of video self-supervised representations across these four downstream factors compared to all previous masked video modeling methods.

4.4 Ablations

In this section, we ablate the hyperparameters of our proposed method SIGMA. For that, we use the Vit-S backbone and train and evaluate (full finetuning) on mini SSv2, a subset with 50% of the data.

Number of prototypes C. In Tab. 6 we report the effect of changing the number of prototypes on the performance of SIGMA. As the number of prototypes increases, the performance also increases and peaks at the moderate number of 4k. Also, the proposed method consistently improves upon VideoMAE [69] regardless of the number of prototypes.

Projection network architecture. We show the effect of changing MLP from a base architecture to a shallower, deeper, and wider version and report the performance in Tab. 7 (left). Shallower architecture in the context of projection networks refers to reducing the number of layers. This results in a network with only 2 layers, each containing 1024 neurons. On the other hand, deeper or wider architectures involve adding more layers (4 instead of 3) or increasing the number of neurons in the same number of layers (2048 instead of 1024). The best performance is achieved when our base MLP projection network is used. Also, the slight differences in performance across various

Table 6. Ablating the number of prototypes for SIGMA with MLP (left) and with DINO (right). A moderate number of prototypes obtain the best performance.

Method	# Prototypes	Top-1	Method	# Prototypes	Top-1
VideoMAE	–	56.9	VideoMAE	–	56.9
SIGMA-MLP	1000	59.7	SIGMA-DINO	1000	62.8
SIGMA-MLP	2000	59.7	SIGMA-DINO	2000	62.8
SIGMA-MLP	3000	59.9	SIGMA-DINO	3000	63.0
SIGMA-MLP	4000	60.3	SIGMA-DINO	4000	63.0
SIGMA-MLP	6000	60.1	SIGMA-DINO	6000	62.9

Table 7. Ablating the projection network architecture (a) and different loss functions (b) Our performance is robust to our φ architecture. For deep targets, a simple L2 loss leads to a trivial solution resulting in low performance, while our SIGMA objective achieves the best results.

(a) Choices for φ (see text for details of architectures)				(b) Different reconstruction losses		
Method	Architecture	Params	Acc	Method	Loss	Acc
(a) VideoMAE	–	–	56.94	(a) VideoMAE	L2	56.9
(b) SIGMA-MLP	base	3M	59.7	(b) SIGMA-MLP	L2	36.4
(d) SIGMA-MLP	shallower	1M	59.2	(d) SIGMA-MLP	Eq.7	60.3
(e) SIGMA-MLP	wider	6M	59.6			
(f) SIGMA-DINO	DINO	22M	63.0			

architectures highlight the robustness of the proposed method.

Loss function. Tab. 7 (right) presents an ablation study of various training objective functions analyzed with our framework. The L2 loss, when combined with deep features as targets, yields significantly lower results in comparison to our clustering-based approach, due to finding trivial shortcut solutions, as discussed in Sec.3.2.

5 Conclusion

In this work, we have proposed a new self-supervised pretraining method SIGMA. To tackle the low semanticity of RGB-based reconstruction targets in current video modeling frameworks, we have introduced a Sinkhorn-clustering-based approach that leads to learnable and more abstract reconstruction targets. These both alleviate the issue of training collapse when simultaneously learning a projection network and can even be combined with pretrained image models, such as DINO. The resulting models outperform the state-of-the-art across a large set of datasets and benchmarks. *Limitations.* Owing to our academic compute budgets, we are not able to run larger architectures, such as ViT-L or ViT-H, which require 64 GPUs [68]. We have, however, demonstrated positive scaling behavior for going from ViT-S to ViT-B.

Acknowledgement This work is financially supported by Qualcomm Technologies Inc., the University of Amsterdam, and the allowance Top consortia for Knowledge and Innovation (TKIs) from the Netherlands Ministry of Economic Affairs and Climate Policy.

References

1. Arnab, A., Deghani, 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 (ICCV) (2021) [2](#)
2. Asano, Y., Patrick, M., Rupprecht, C., Vedaldi, A.: Labelling unlabelled videos from scratch with multi-modal self-supervision. In: Advances in Neural Information Processing Systems (2020) [6](#)
3. Asano, Y.M., Rupprecht, C., Vedaldi, A.: Self-labelling via simultaneous clustering and representation learning. In: ICLR (2020) [3](#)
4. Assran, M., Duval, Q., Misra, I., Bojanowski, P., Vincent, P., Rabbat, M., LeCun, Y., Ballas, N.: Self-supervised learning from images with a joint-embedding predictive architecture. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15619–15629 (2023) [3](#)
5. Bardes, A., Ponce, J., LeCun, Y.: Mc-jepa: A joint-embedding predictive architecture for self-supervised learning of motion and content features (2023) [2, 7](#)
6. 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) [3](#)
7. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? ArXiv [abs/2102.05095](#) (2021) [9, 10](#)
8. Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: Proceedings of the European conference on computer vision (ECCV). pp. 132–149 (2018) [3](#)
9. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments (2020) [6](#)
10. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9650–9660 (2021) [7](#)
11. Chen, L., Wu, P., Chitta, K., Jaeger, B., Geiger, A., Li, H.: End-to-end autonomous driving: Challenges and frontiers. arXiv preprint arXiv:2306.16927 (2023) [1](#)
12. Chen, P., Huang, D., He, D., Long, X., Zeng, R., Wen, S., Tan, M., Gan, C.: Rspnet: Relative speed perception for unsupervised video representation learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 1045–1053 (2021) [3](#)
13. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020) [3](#)
14. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transportation distances (2013) [6](#)
15. Dave, I., Gupta, R., Rizve, M.N., Shah, M.: Tclr: Temporal contrastive learning for video representation. Computer Vision and Image Understanding **219**, 103406 (2022) [3](#)
16. Deghani, M., Djolonga, J., Mustafa, B., Padlewski, P., Heek, J., Gilmer, J., Steiner, A.P., Caron, M., Geirhos, R., Alabdulmohsin, I., et al.: Scaling vision transformers to 22 billion parameters. In: International Conference on Machine Learning. pp. 7480–7512. PMLR (2023) [2](#)
17. 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 (2009) [8](#)
18. Dorkenwald, M., Xiao, F., Brattoli, B., Tighe, J., Modolo, D.: Scvrl: Shuffled contrastive video representation learning (2022) [2, 3, 7](#)

19. 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) [2](#), [3](#)
20. Duan, H., Zhao, N., Chen, K., Lin, D.: Transrank: Self-supervised video representation learning via ranking-based transformation recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3000–3010 (2022) [3](#)
21. Duan, J., Yu, S., Tan, H.L., Zhu, H., Tan, C.: A survey of embodied ai: From simulators to research tasks. IEEE Transactions on Emerging Topics in Computational Intelligence **6**(2), 230–244 (2022) [1](#)
22. Fan, D., Wang, J., Liao, S., Zhu, Y., Bhat, V., Santos-Villalobos, H.J., MV, R., Li, X.: Motion-guided masking for spatiotemporal representation learning. In: ICCV. pp. 5596–5606. IEEE (2023) [2](#), [4](#), [8](#), [9](#), [10](#), [12](#)
23. Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., Feichtenhofer, C.: Multiscale vision transformers. In: ICCV. pp. 6804–6815. IEEE (2021) [2](#), [9](#), [10](#)
24. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: ICCV (2018) [9](#), [10](#)
25. Feichtenhofer, C., Li, Y., He, K., et al.: Masked autoencoders as spatiotemporal learners. Advances in neural information processing systems **35**, 35946–35958 (2022) [2](#), [3](#), [4](#), [5](#)
26. 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) [1](#), [3](#)
27. Finn, C., Levine, S.: Deep visual foresight for planning robot motion. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). pp. 2786–2793. IEEE (2017) [1](#)
28. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. arXiv preprint arXiv:1803.07728 (2018) [3](#)
29. Girdhar, R., El-Nouby, A., Singh, M., Alwala, K.V., Joulin, A., Misra, I.: Omnimae: Single model masked pretraining on images and videos (2023) [9](#), [10](#)
30. 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) [7](#)
31. Ha, D., Schmidhuber, J.: World models. arXiv preprint arXiv:1803.10122 (2018) [1](#)
32. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022) [3](#)
33. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020) [3](#)
34. Huang, B., Zhao, Z., Zhang, G., Qiao, Y., Wang, L.: Mgm: Motion guided masking for video masked autoencoding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13493–13504 (2023) [4](#), [11](#), [12](#)
35. Huang, B., Zhao, Z., Zhang, G., Qiao, Y., Wang, L.: MGMAE: motion guided masking for video masked autoencoding. In: ICCV. pp. 13447–13458. IEEE (2023) [8](#), [9](#), [10](#), [12](#), [13](#)
36. Janai, J., Güney, F., Behl, A., Geiger, A., et al.: Computer vision for autonomous vehicles: Problems, datasets and state of the art. Foundations and Trends® in Computer Graphics and Vision **12**(1–3), 1–308 (2020) [1](#)
37. 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) [3](#)

38. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009) [8](#)
39. 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) [8](#)
40. Kuhn, H.W.: The hungarian method for the assignment problem. *Naval research logistics quarterly* **2**(1-2), 83–97 (1955) [10](#)
41. Kuutti, S., Bowden, R., Jin, Y., Barber, P., Fallah, S.: A survey of deep learning applications to autonomous vehicle control. *IEEE Transactions on Intelligent Transportation Systems* **22**(2), 712–733 (2020) [1](#)
42. Lee, H.Y., Huang, J.B., Singh, M., Yang, M.H.: Unsupervised representation learning by sorting sequences. In: Proceedings of the IEEE international conference on computer vision. pp. 667–676 (2017) [3](#)
43. Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: CVPR. pp. 3192–3201. IEEE (2022) [9](#), [10](#)
44. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization (2019) [7](#)
45. Lu, C., Jin, X., Huang, Z., Hou, Q., Cheng, M., Feng, J.: CMAE-V: contrastive masked autoencoders for video action recognition. *CoRR* **abs/2301.06018** (2023) [2](#), [4](#), [8](#), [9](#), [10](#)
46. Matsuo, Y., LeCun, Y., Sahani, M., Precup, D., Silver, D., Sugiyama, M., Uchibe, E., Morimoto, J.: Deep learning, reinforcement learning, and world models. *Neural Networks* **152**, 267–275 (2022) [1](#)
47. Misra, I., Zitnick, C.L., Hebert, M.: Shuffle and learn: unsupervised learning using temporal order verification. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14. pp. 527–544. Springer (2016) [3](#)
48. Nguyen, D.K., Aggarwal, V., Li, Y., Oswald, M.R., Kirillov, A., Snoek, C.G.M., Chen, X.: R-mae: Regions meet masked autoencoders. In: ICLR (2024) [3](#)
49. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: European conference on computer vision. pp. 69–84. Springer (2016) [3](#)
50. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018) [3](#)
51. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2536–2544 (2016) [3](#)
52. Patrick, M., Asano, Y., Kuznetsova, P., Fong, R., Henriques, J.F., Zweig, G., Vedaldi, A.: Multi-modal self-supervision from generalized data transformations (2020) [2](#), [7](#)
53. Photutorial.com: Key youtube video statistics in 2024 (2024) [1](#)
54. Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675* (2017) [10](#)
55. Qian, R., Meng, T., Gong, B., Yang, M.H., Wang, H., Belongie, S., Cui, Y.: Spatiotemporal contrastive video representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6964–6974 (2021) [2](#), [3](#), [7](#)
56. Ranasinghe, K., Naseer, M., Khan, S., Khan, F.S., Ryoo, M.S.: Self-supervised video transformer. In: CVPR. pp. 2864–2874. IEEE (2022) [7](#), [12](#), [13](#)
57. Richardson, I.E.: Video codec design: developing image and video compression systems. John Wiley & Sons (2002) [4](#)
58. Salehi, M., Gavves, E., Snoek, C.G., Asano, Y.M.: Time does tell: Self-supervised time-tuning of dense image representations. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16536–16547 (2023) [3](#), [10](#), [12](#)

59. Shao, D., Zhao, Y., Dai, B., Lin, D.: Finegym: A hierarchical video dataset for fine-grained action understanding. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2616–2625 (2020) [12](#)
60. Sigurdsson, G.A., Varol, G., Wang, X., Farhadi, A., Laptev, I., Gupta, A.: Hollywood in homes: Crowdsourcing data collection for activity understanding. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14. pp. 510–526. Springer (2016) [12](#)
61. Singh, M., Duval, Q., Alwala, K.V., Fan, H., Aggarwal, V., Adcock, A., Joulin, A., Dollár, P., Feichtenhofer, C., Girshick, R., Girdhar, R., Misra, I.: The effectiveness of mae pre-training for billion-scale pretraining. In: ICCV (2023) [2](#)
62. 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) [8](#)
63. 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: CVPR. pp. 2235–2245. IEEE (2023) [2, 4, 5, 8, 9](#)
64. Tao, L., Wang, X., Yamasaki, T.: Self-supervised video representation learning using inter-intra contrastive framework. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 2193–2201 (2020) [3](#)
65. Thoker, F.M., Doughty, H., Bagad, P., Snoek, C.G.M.: How severe is benchmark-sensitivity in video self-supervised learning? In: European Conference on Computer Vision (ECCV) (2022) [3, 7, 12, 13](#)
66. Thoker, F.M., Doughty, H., Snoek, C.: Tubelet-contrastive self-supervision for video-efficient generalization (2023) [3, 12](#)
67. Thoker, F.M., Doughty, H., Snoek, C.G.M.: Tubelet-contrastive self-supervision for video-efficient generalization. In: ICCV. pp. 13766–13777. IEEE (2023) [3, 7, 13](#)
68. Tong, Z., Song, Y., Wang, J., Wang, L.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In: NeurIPS (2022) [2, 4, 7, 8, 9, 10, 12, 14](#)
69. Tong, Z., Song, Y., Wang, J., Wang, L.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. Advances in neural information processing systems **35**, 10078–10093 (2022) [3, 11, 13](#)
70. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Canton-Ferrer, C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P.S., Lachaux, M., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E.M., Subramanian, R., Tan, X.E., Tang, B., Taylor, R., Williams, A., Kuan, J.X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., Scialom, T.: Llama 2: Open foundation and fine-tuned chat models. CoRR **abs/2307.09288** (2023) [1](#)
71. 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) [3](#)
72. Wang, J., Gao, Y., Li, K., Hu, J., Jiang, X., Guo, X., Ji, R., Sun, X.: Enhancing unsupervised video representation learning by decoupling the scene and the motion. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 10129–10137 (2021) [3](#)
73. 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 (CVPR). pp. 14549–14560 (June 2023) [2, 4](#)

74. 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) [3](#), [4](#)
75. 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) [4](#)
76. Wang, R., Chen, D., Wu, Z., Chen, Y., Dai, X., Liu, M., Yuan, L., Jiang, Y.: Masked video distillation: Rethinking masked feature modeling for self-supervised video representation learning. In: CVPR. pp. 6312–6322. IEEE (2023) [8](#), [12](#), [13](#)
77. 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) [3](#)
78. Wellhausen, L., Dosovitskiy, A., Ranftl, R., Walas, K., Cadena, C., Hutter, M.: Where should i walk? predicting terrain properties from images via self-supervised learning. *IEEE Robotics and Automation Letters* **4**(2), 1509–1516 (2019) [1](#)
79. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3733–3742 (2018) [3](#)
80. Xie, Z., Zhang, Z., Cao, Y., Lin, Y., Bao, J., Yao, Z., Dai, Q., Hu, H.: Simmim: A simple framework for masked image modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9653–9663 (2022) [3](#)
81. 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) [3](#)
82. Xu, N., Yang, L., Fan, Y., Yue, D., Liang, Y., Yang, J., Huang, T.: Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327* (2018) [10](#)
83. 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) [5](#)
84. Yao, T., Zhang, Y., Qiu, Z., Pan, Y., Mei, T.: Seco: Exploring sequence supervision for unsupervised representation learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 10656–10664 (2021) [3](#)
85. 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) [1](#), [3](#)
86. Zhang, H., Xu, X., Han, G., He, S.: Context-aware and scale-insensitive temporal repetition counting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [12](#)
87. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14. pp. 649–666. Springer (2016) [3](#)