

Spatio-Temporal Crop Aggregation for Video Representation Learning

Sepehr Sameni
 University of Bern
 sepehr.sameni@unibe.ch

Simon Jenni
 Adobe Research
 jenni@adobe.com

Paolo Favaro
 University of Bern
 paolo.favaro@unibe.ch

Abstract

We propose *Spatio-temporal Crop Aggregation for video representation LEarning (SCALE)*, a novel method that enjoys high scalability at both training and inference time. Our model builds long-range video features by learning from sets of video clip-level features extracted with a pre-trained backbone. To train the model, we propose a self-supervised objective consisting of masked clip feature predictions. We apply sparsity to both the input, by extracting a random set of video clips, and to the loss function, by only reconstructing the sparse inputs. Moreover, we use dimensionality reduction by working in the latent space of a pre-trained backbone applied to single video clips. These techniques make our method not only extremely efficient to train but also highly effective in transfer learning. We demonstrate that our video representation yields state-of-the-art performance with linear, nonlinear, and k -NN probing on common action classification and video understanding datasets.

1. Introduction

Videos provide rich and detailed information about objects and their activities. Their analysis is, however, made challenging not only by the difference in the information provided across space and time but also by the high dimensionality of the data [45, 29, 70]. While computational resources are expected to scale over time, so is the demand for higher data resolution (both in space and time) and also the need for processing data with even more dimensions, such as videos of volumetric data [10]. Therefore, it is of paramount importance to explore methods that drastically reduce the computational requirements to process videos. Moreover, the annotation of videos is an extremely costly and time-consuming burden that makes the use of models pre-trained in a self-supervised manner essential [44].

Self-supervised learning (SSL) is a very popular technique to reduce the need for annotation because it can build useful representations from unlabeled data through an artificial goal, also called pseudo- or pretext-task. These repre-

sentations can either be evaluated through K -nearest neighbor or linear probing [71, 72] or through fine-tuning (i.e., as the initialization parameters of the trained model) [25, 54, 35] on a downstream task, where only a small labeled dataset is available. More remarkably, SSL pre-trained models can outperform models that were pre-trained on an annotated dataset [53, 4].

SSL methods for video representation learning present fundamental scalability challenges [54, 45, 29, 65]. A first major challenge is that training models from scratch for any new pseudo-task is not sustainable. A more viable setting is one where the processing of videos is broken down into multiple steps. First, one pre-processes video clips (spatio-temporal crops) only once via some pre-trained general-purpose model (e.g., trained via self-supervised learning [18, 23, 26, 44, 7, 37, 41]). In general, the idea is that this (compressed) local representation is stored and used later for other training purposes or retrieval. Then, one can train more efficiently other super-features that aggregate information from these local representations. Within this setting, we address two questions: 1) Can one further improve the performance of pre-computed video features by training a model on top of them? 2) Can such training be made computationally scalable? In this paper, we show that both questions have a positive answer. The answer exploits four fronts:

1. **Input Sparsity:** Sparsity in the input to the model [54, 25, 3, 19, 1] drastically reduces computational load and memory requirements, leveraging information redundancy in images and videos [16]. Inspired by prior work, we extract a sparse set of clips from a video (these are spatio-temporal video crops and not image patches or video tubelets); each clip is fed separately to a neural network so that we obtain a set of video clip representations.
2. **Output Sparsity:** Another way to reduce the computational cost is to use a sparse reconstruction output instead of a dense one [71, 51, 59, 5]. This is not just a reduction of the number of terms in the loss function, but also in the number of actual outputs of

the model, and thus also a reduction of the computations needed to obtain them. This is in contrast to MAE-based SSL methods for vision, where the proposed pseudo-tasks are based on the reconstruction of the whole input [25, 54, 19, 16], and even if the loss uses a subset of the tokens (the masked ones) for the loss calculation, the remaining tokens are still part of the decoder’s output and computational graph.

3. **Dimensionality Reduction:** Inspired by prior work [71, 41, 13], instead of directly processing the raw input data, we work in the latent space. This allows us to further reduce the dimensionality of both input and output data.
4. **Use of a Pre-trained Backbone:** To reduce training time and further speed up the processing per iteration, we exploit SSL pre-trained models. These pre-trained backbones already learn very strong short-term spatio-temporal features, and our approach is a way to extract longer-term features (see LVU results 2) by building a video representation on top of a set of pre-trained features (one for each video clip).

To integrate all these components, we propose a novel SSL method that we call Spatio-temporal Crop Aggregation for video representation LEarning (SCALE). Given an input video, SCALE extracts a random set of clips and produces an embedding for each clip through a pre-trained backbone, which is kept frozen. These initial embeddings are then augmented in two ways: 1) each one is refined into a more discriminative feature, and 2) the set of all embeddings is summarized in a global feature. These global features can learn long-term correlations in the whole video by aggregating the short-term information in each clip embedding. The combination of the initial embeddings with their refinement and the global feature is then used in an ensemble for applications on new downstream tasks.

To train SCALE, we introduce two novel pseudo-tasks, which aim to improve the discriminability of the embeddings of each video clip as well as obtain a global representation of the video. One task, which we call Masked Clip Modeling (MCM), is the reconstruction of a video clip embedding as in masked autoencoders [54]. Masked embeddings are combined with positional encodings so that the model can (spatio-temporally) relate the missing input embeddings to the other available embeddings (similarly to BERT in Natural Language Processing [12]). A second task is to train the model to output a global feature token (which we refer to as CLS, just for consistency with previous works [7, 12, 14, 55, 9]) for a set of clips via contrastive learning, so that the global feature is invariant to the chosen set of clips from the same video, but can discriminate summary features of clips from other videos. Both tasks are trained via contrastive losses.

SCALE consistently improves upon pretrained features, with more evident gains considering the computational cost of other methods. For instance, SCALE outperforms VideoMAE [54] by 0.3% (1600 epochs checkpoint – fine-tuning) and ρ BYOL [18] by 0.6% (800 epochs checkpoint – linear probe) on Kinetics400 with 64 V100 GPUs in about **5 mins**. In contrast, VideoMAE takes **27.7 hours** to gain 0.5% (800 to 1600 epochs checkpoint), and ρ BYOL needs **48 hours** to gain 0.4% (200 to 400 epochs checkpoint) with the same resources.

To summarize our contributions, we propose SCALE, a novel and highly scalable video representation method that

- is trained via novel pseudo-tasks on sets of video clips (in contrast to existing methods that work only with pairs of clips at a time [44, 18, 45]);
- results in video feature representations with a significant performance improvement in k -NN (retrieval), linear, and nonlinear probing across a wide range of datasets for action classification (UCF [48], HMDB [34], SSv2 [22], Kinetics400 [33]) and long-form video understanding (LVU [63]);
- achieves consistent transfer learning performance improvement across diverse SotA pre-trained backbones (architectures, scale, and pre-training tasks). For example, our nonlinear probed model even outperforms fully fine-tuned SVT [44] on HMDB [34].

2. Related Work

Our approach relates to many prior works on (self-supervised) representation learning. In particular, SCALE relates to SSL approaches on videos, methods that rely on multiple views of the data, and predictive methods, where part of the data is predicted from another part.

Self-Supervised Video Representation Learning. Early SSL approaches on video were based on pseudo tasks, *e.g.*, the recognition of transformations of video frame sequences [38, 6, 62]. More recently, popular methods developed on images have been successfully translated to video, *e.g.*, contrastive methods [18, 23, 26], clustering-based methods [44, 7], or predictive approaches [37, 41]. Often, these approaches are tailored to video data by including additional learning signals, *e.g.*, by combining contrastive methods with temporal constraints [11] and pretext tasks [32], or by incorporating audio [39], or optical flow [24]. These approaches typically learn representations with limited temporal extent, which can serve as backbones for our approach.

Learning from Multiple Views. Several methods have exploited multiple views (*e.g.*, generated through space-time cropping) to represent and learn from videos. For example, many SSL approaches rely on multiple views of the

data, *e.g.*, in contrastive formulations [18, 44, 42], or predictive learning [45], where invariance to views is the goal. Other approaches aim to learn from the relation between two views, *e.g.*, by predicting overlap [69], the relative distance [49], or cross-view feature prediction [68, 52] and reconstruction [40]. Besides exploiting multiple views for SSL, some works also propose general multi-view video models, *e.g.*, by capturing and fusing features at different spatio-temporal resolutions [67], by aggregating information over longer time-spans [64, 47, 65, 57], or by selecting important frames [21]. These approaches, however, are learned end-to-end. In contrast, we propose a more scalable approach by learning a global video representation of pre-trained features extracted over multiple crops using self-supervision.

Predictive Learning. One of our proposed SSL objectives is a prediction of space-time crop features given other video crop features. This approach is similar to masked token prediction as in BERT [12] and relates to several other methods in the literature. Masked input reconstruction methods have recently become popular on images [25] and successfully translated to video [54, 16]. Other approaches formulate masked prediction tasks in the learned feature space [71, 52, 13] or some fixed latent space [51]. Another line of work considers directional predictions (*e.g.*, into the future) often formulated via contrastive predictive coding [41] applied to video [37, 36, 50, 66, 20]. These masked prediction tasks are typically formulated on a fixed grid (*e.g.*, at the level of patches or frames). In contrast, our formulation is continuous, predicting features of arbitrarily sampled space-time crops.

3. Scalable Video Representation Learning

To describe SCALE, we first define some basic notation and functions, including a general contrastive loss notation that we use for all training losses.

3.1. Notation

We use lower-case letters (*e.g.*, z) to denote generic vectors and capital letters (*e.g.*, Z) to denote their sets. The expression $a \odot b$ denotes the concatenation of a and b . We also avoid indicating the parameters of networks (usually denoted by θ) if their presence and role are clear from the context. Throughout the description of the method we refer to the training of neural networks and therefore at each iteration of the training we sample a minibatch of videos. All the equations in the next sections are written for a single video in the minibatch. Although we do not explicitly indicate it, all the contrastive losses also use the other videos in the minibatch as negatives.

3.2. Contrastive Loss

InfoNCE is a powerful method for representation learning [41] that can be used to maximize the mutual information between two variables. Because we use this loss between different variables throughout our method, we introduce here a unified notation. Let the paired sets A and B have N elements each, $A = \{a^i\}_{i=1}^N$ and $B = \{b^i\}_{i=1}^N$, where a^i are vectors of dimension d_A and b^i are vectors of dimension d_B . We also introduce two Multi Layer Perceptrons (MLP), parameterized with θ_A and θ_B , to project these vectors onto a common space of dimension d . After feeding the elements a^i and b^i to the MLPs and normalizing them, we obtain

$$\tilde{a}^i = \frac{\text{MLP}_{\theta_A}(a^i)}{\|\text{MLP}_{\theta_A}(a^i)\|} \quad \text{and} \quad \tilde{b}^i = \frac{\text{MLP}_{\theta_B}(b^i)}{\|\text{MLP}_{\theta_B}(b^i)\|}, \quad (1)$$

where $\|\cdot\|$ denotes the L_2 norm. We define the per-element loss based on the relative similarity of \tilde{a}^i and \tilde{b}^i , and by using a temperature τ

$$\tilde{\ell}^i(A, B, \theta_A, \theta_B) = -\log \frac{\exp\left(\frac{\tilde{a}^i \cdot \tilde{b}^i}{\tau}\right)}{\sum_{j=1}^N \exp\left(\frac{\tilde{a}^i \cdot \tilde{b}^j}{\tau}\right)}. \quad (2)$$

We then make the loss symmetric [43] by using

$$\ell^i(A, B, \theta_A, \theta_B) = \tilde{\ell}^i(A, B) + \tilde{\ell}^i(B, A). \quad (3)$$

Finally, we define the contrastive loss $\mathcal{L}_{\text{ctr}}(A, B, \theta_A, \theta_B)$ as the mean of ℓ^i

$$\mathcal{L}_{\text{ctr}}(A, B, \theta_A, \theta_B) = \frac{1}{N} \sum_{i=1}^N \ell^i(A, B, \theta_A, \theta_B). \quad (4)$$

As mentioned earlier on, for simplicity, in the rest of the paper we will not indicate the parameters of the MLPs, and simply write $\mathcal{L}_{\text{ctr}}(A, B)$ or $\ell^i(A, B)$.

3.3. Training SCALE

In our method, we integrate 4 principles to drastically reduce the computational complexity of learning a video representation: Input sparsity, output sparsity, dimensionality reduction, and use of a pre-trained backbone. Moreover, we introduce two pseudo-tasks to train the model. One task is based on the (contrastive) reconstruction of a masked video clip given some context video clips. The second task is to build a global representation that is (contrastively) invariant to the set of input sampled video clips. The overall training scheme of SCALE is illustrated in Figure 1.

Input sparsity. As a first step, rather than processing a whole video, we collect a sparse set of short video clips from the same video. Given a video $V \in \mathbb{R}^{H \times W \times T \times 3}$,

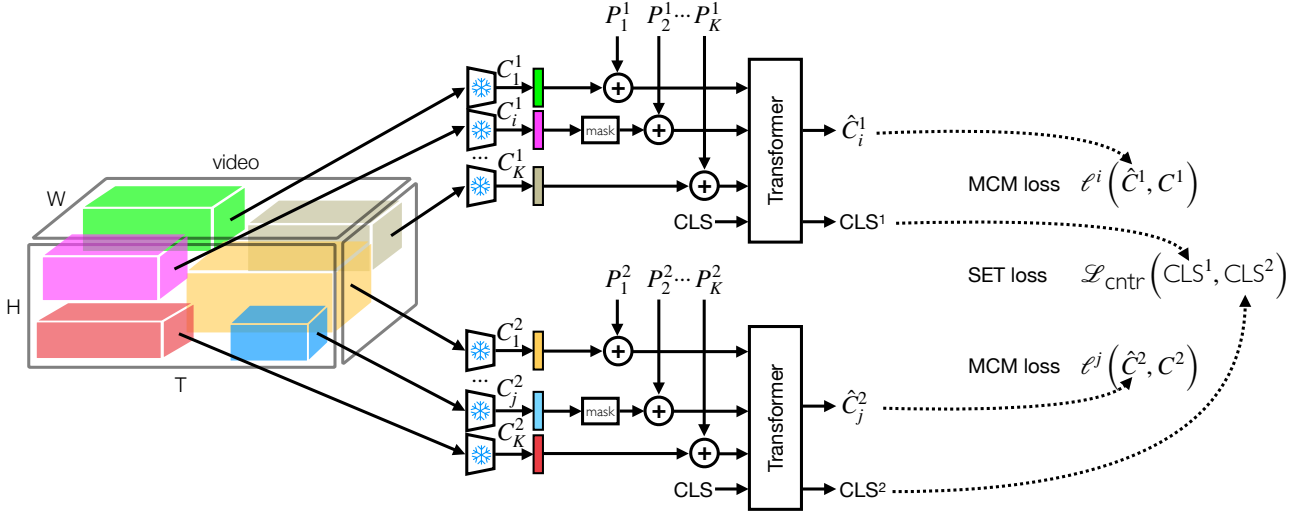


Figure 1. Video Representation Learning with SCALE. For each video, SCALE extracts two sets of video clips V_1^1, \dots, V_K^1 and V_1^2, \dots, V_K^2 . Each video clip is processed separately through a frozen backbone E_* and results in encoded video clips C_1^1, \dots, C_K^1 and C_1^2, \dots, C_K^2 . Then, a random set of encodings in each set is masked and reconstructed at the output of the predictor network (a transformer) (ℓ^i). The predictor network is also fed a class token CLS. The corresponding output token encodes a summary CLS^m of the m -th set of video clips. The objective for these summary tokens is to be similar only when encoding video clips from the same video (\mathcal{L}_{SET}).

where H , W and T are the height, width, and duration (in frames) of the video, we sample $2 \times K$ clips. We divide the clips into two sets randomly. Each clip in the first set $V_i^1 \in \mathbb{R}^{H_i^1 \times W_i^1 \times T_i^1 \times 3}$, with $i = 1, \dots, K$, is obtained at the spatio-temporal location X_i^1, Y_i^1, Q_i^1 with different data augmentations and dimensions H_i^1, W_i^1 and T_i^1 . Similarly, we denote the second set of clips V_i^2 , for $i = 1, \dots, K$. We also normalize their coordinates relative to the video dimensions and embed them onto a feature vector P_i^j by feeding them to a learnable MLP. We denote these embeddings

$$P_i^j = \text{MLP} \left(\left[\frac{X_i^j}{H}, \frac{Y_i^j}{W}, \frac{Q_i^j}{T}, \frac{X_i^j + H_i^j}{H}, \frac{Y_i^j + W_i^j}{W}, \frac{Q_i^j + T_i^j}{T} \right]^\top \right), \quad (5)$$

where $j = 1, 2$ and $i = 1, \dots, K$.

Dimensionality reduction and pre-trained backbone. To reduce the dimensionality of each video clip, we feed them independently to a frozen encoder E_* , to obtain the encodings $C_i^j = E_*(V_i^j)$, where $j = 1, 2$ and $i = 1, \dots, K$. Our framework is encoder-agnostic and thus can work with encoders obtained through different training schemes (supervised, contrastive, or autoencoder). In addition to reducing the dimensionality, we make the training even more scalable by using pre-trained and frozen encoders. Note, however, that if performance is the main goal, it is possible to also train a sparse backbone end to end with multiple clips. Thanks to *token dropping*, one can drop up to 95% of the tokens [19] and still build a good representation.

Output sparsity. As a self-supervised signal for our video

representation learning, we use a (contrastive) reconstruction loss. To reduce the computational cost, instead of predicting the features for the whole video [59, 19, 54, 16] (asymmetric decoding), we only reconstruct a *sparse* set of masked clips. Our reconstruction objective is based on the observation that video signals carry a lot of redundancy. Hence, we introduce a model, the *predictor network*, to predict masked video clip embeddings given the other video clip embeddings (the context). We follow the general approach of BERT [12] but implement the predictor network as a masked autoencoder, where the reconstruction is based on a contrastive loss. The loss is applied only to a sparse set $M^1 \subset \{1, \dots, K\}$ of masked video clips. These clips are replaced by a learned MSK token. All embeddings C_i^1 , including the masked ones, are added to their corresponding position encoding P_i^1 and are then fed to the predictor network. We also include an additional learnable CLS token as input for the predictor network, which will be used for tasks with multiple video clips. We denote the outputs of the predictor network as \hat{C}_i^1 for the tokens corresponding to C_i^1 , and as CLS^1 for the token corresponding to CLS. Similarly, we feed as inputs separately from the previous set all the video clips C_i^2 with their corresponding positional embeddings P_i^2 , and the same CLS token, and obtain \hat{C}_i^2 and CLS^2 respectively (see Figure 1 for a visual depiction of these processing steps).

Contrastive reconstruction. Modeling all the details of a masked clip, even in the latent space and even given the redundancy in videos, is a demanding task. Rather than

increasing the capacity of our model, since we are aiming for scalability, we keep our predictor network a (relatively) shallow network and use contrastive learning [41]. With contrastive learning, the predicted representation of the masked tokens should only be closer to the original unmasked clip representation (after an MLP projection) than from all other clips from the same video and the rest of the minibatch. Note that the rest of the clips in the same video act as hard negatives in contrastive learning [46]. Also, since we are using a frozen backbone, we can afford to use large minibatch sizes, which is known to be beneficial for contrastive learning [9]. We call this contrastive reconstruction loss the Masked Clip Modeling (MCM) loss

$$\mathcal{L}_{\text{MCM}} = \sum_{i \in M^1} \ell^i(\hat{C}^1, C^1) + \sum_{j \in M^2} \ell^j(\hat{C}^2, C^2). \quad (6)$$

Multiple video clips loss. The predictor network outputs features for each video clip that are highly discriminative. This task is similar to that of a masked autoencoder [25] and gives you an enhanced per-clip representation. For many video tasks, we need a global representation for the whole video; for that, we introduce an additional pseudo-task that captures a more global representation of a set of video clips. Our task takes inspiration from contrastive learning methods used in SSL, which yield representations that perform well with linear probing [8]. The loss aims to make the CLS^1 and CLS^2 tokens returned from the predictor network more similar (recall that these two tokens are obtained from two separate groups of video clips extracted from the same video) than to other class tokens from other videos within the minibatch

$$\mathcal{L}_{\text{SET}} = \mathcal{L}_{\text{ctr}}(\text{CLS}^1, \text{CLS}^2). \quad (7)$$

In addition to our contrastive loss (InfoNCE), one can use clustering [7] or regression [23] losses. We choose InfoNCE for simplicity and for better compatibility between the losses. As the overall loss, we use the sum of both loss terms (without any weights)

$$\mathcal{L} = \mathcal{L}_{\text{MCM}} + \mathcal{L}_{\text{SET}}, \quad (8)$$

4. Experiments

We evaluate SCALE on several commonly used action classification datasets for video representation learning. As our performance metric, we primarily use linear probing and nonlinear probing [27]. For the smaller datasets, we also use k -NN classifiers (which are training-free) and demonstrate that the proposed method improves upon both unsupervised and supervised backbones.

4.1. Experimental Setup and Protocols

Computational Efficiency: In Table 1, we show estimates of the maximum batch sizes and the training throughput for

Table 1. **Computational Resources.** GPU VRAM usage (GB) and max training speed (using one 3090 GPU) of ViT_B with varying batch sizes and SSL tasks for videos. MoCo_{Sparse}^{V3} is akin to MSN [3]. OOM indicates Out Of Memory error.

Batch Size	8	19	57	2048	Samples/s
MoCo ^{V3}	23.47	OOM	OOM	OOM	8.66
VMAE	11.42	22.80	OOM	OOM	33.54
MoCo _{Sparse} ^{V3}	5.49	9.27	23.62	OOM	28.83
SCALE	1.21	1.23	1.58	19.52	9224.65

different methods trained with the same computational and memory resources. As can be seen, SCALE is orders of magnitude more efficient than other SotA methods. Also during multi-crop inference, our method only results in an FLOP increase of approximately 0.001%.

Datasets: Following prior work [44, 18, 45] we use Kinetics-400 [33], UCF-101 [48] (split 1), HMDB-51 [34] (split 1), and Something-Something v2 (SSv2) [22] to train and evaluate our models. We also use the LVU benchmark [63] to showcase our long-form video understanding capabilities. Note that almost 35% of LVU videos are not available to download from YouTube anymore; thus, our results are not directly comparable with prior methods.

Pretrained backbones: We use the pretrained checkpoints of ρ BYOL [18], SVT [44], and three variants of VideoMAE [54] (base(B), large(L), and fine-tuned base(FT)). We choose ρ BYOL for their excellent linear performance, SVT for the usage of ViT [14], and VMAE for showing 1) the applicability of our proposed method to MAE models, 2) the scalability of our method to larger models, and 3) possibility of using supervisedly fine-tuned models as our backbone. All the models are self-supervisedly pretrained on Kinetics-400, except the fine-tuned VMAE base that was also supervisedly finetuned on Kinetics-400. We also used a backbone pretrained and fine-tuned on SSv2 (VMAE_{SSv2}^B) for the SSv2 experiment to show the universality of SCALE with respect to the pretraining dataset.

Self-supervised Training: For training data, we extract 16 clips of 16 frames from each video per dataset and save their feature encodings to disk. We use PySlowFast’s common data augmentations for that [15]. For evaluation, we follow the 5×3 scheme [17] (uniformly sampling 5 clips from a video along its temporal axis and then taking 3 spatial crops) and extract 15 clips from each video (except for SSv2, where we extract 2×3 clips [58]). As the architecture for the predictor network, we use an encoder-only Transformer [56] and a three-layer MLP (without batch normalization [30]) for the contrastive heads. Unless stated otherwise, we train our models for 500 epochs (for example, training with VMAE^B on SSv2 takes 137 minutes with one 3090 GPU) with a batch size of 512 and use all 16 clips.

Evaluation: Since our focus is on efficient and scalable

Table 2. **Long-Form Video Understanding Results.** Linear and nonlinear probing accuracies on LVU [63] classification tasks. SCALE shows significant performance improvement compared to baselines, indicating its ability to capture long-form video features. The first two rows are greyed out as direct comparison is not possible due to the unavailability of the full dataset for download.

		Relation	Speak	Scene	Director	Genre	Writer	Year
SlowFast+NL	[17, 60]	52.40	35.80	54.70	44.90	53.00	36.30	52.50
ViS4mer	[31]	57.14	40.79	67.44	62.61	54.71	48.80	44.75
SVT	Linear	64.70	35.77	62.33	37.27	54.35	52.12	28.46
	SCALE _{linear}	73.52 (+8.82)	40.65 (+4.88)	68.83 (+6.50)	46.36 (+9.09)	57.94 (+3.59)	56.38 (+4.26)	39.23 (+10.77)
	MLP	67.64	39.02	66.23	45.45	56.92	57.44	36.15
	Transformer	70.58	40.65	68.83	47.27	57.17	58.51	36.92
	SCALE _{ft}	76.47 (+5.89)	42.27 (+1.62)	74.02 (+5.19)	49.09 (+1.82)	58.97 (+1.80)	62.76 (+4.25)	39.23 (+2.31)
ρ BYOL	Linear	52.94	38.21	53.24	36.36	51.79	57.44	33.84
	SCALE _{linear}	67.64 (+14.70)	43.08 (+4.87)	66.23 (+12.99)	44.54 (+8.18)	53.33 (+1.54)	60.63 (+3.19)	40.00 (+6.16)
	MLP	62.35	41.46	62.42	47.27	52.56	59.57	40.00
	Transformer	65.09	44.71	66.23	50.90	53.07	61.70	43.84
	SCALE _{ft}	67.64 (+2.55)	45.52 (+0.81)	71.42 (+5.19)	51.81 (+0.91)	55.72 (+2.65)	65.95 (+4.25)	46.92 (+3.08)

video classification, we always freeze the backbones in our evaluation (as in our self-supervised pretraining) and either train a linear classifier [44, 18] or fine-tune the predictor network (the transformer) with an additional linear head. Therefore, when we refer to fine-tuning (ft), we *only* adapt the nonlinear head (e.g., predictor network) but *not the backbone*. We apply a grid search for the hyper-parameters of the heads covering learning rate, weight decay, batch size, and optimizer type. Similar to MAE [25], we found that applying a batch normalization layer [30] without affine transformations is beneficial for VideoMAE models. As the linear baseline, we consider the well-established ensembling approach, *i.e.*, we average the softmax predictions of the 15 clips (6 for SSv2) to obtain the final prediction. For models that process multiple clips at once (like ours), we likewise apply a linear softmax head on the concatenation of the individual clip features and the [SET] token before averaging to obtain the final prediction. For the smaller datasets, we also use k -NN classification, where, similar to DINO [7], we always use $k = 20$ and work with l2 normalized representations.

Nonlinear baselines: As SCALE is a nonlinear model, we consider an MLP on top of the frozen backbone as a nonlinear baseline. As a further baseline and to illustrate the effect of our self-supervised pre-training, we consider a Transformer trained on all the clip representations. This Transformer uses the exact same architecture as SCALE, and only differs in the initialization: in the case of SCALE we start from our proposed SSL pre-trained weights instead of random initialization.

4.2. Results

LVU: One of the benefits of SCALE is that it can be used to process long videos, even though the backbones were trained on short videos only. To demonstrate the ability to capture long-form video features, we evaluated SCALE on LVU [63], a benchmark that involves seven classifica-

Table 3. **SSv2 Results.** Linear and nonlinear probing accuracies on SSv2 [22]. We see that both SCALE_{linear} and SCALE_{ft} outperform other methods and improve the classification accuracies by a large margin. We also see that SCALE_{ft}, with its better initialization, always outperforms the Transformer. $VMAE_{SSv2}^B$ was pretrained and fine-tuned on SSv2.

	SVT	ρ BYOL	$VMAE^B$	$VMAE^L$	$VMAE_{ft}^B$	$VMAE_{SSv2}^B$
Linear	20.30	25.30	18.31	27.94	28.90	70.53
SCALE _{linear}	25.26	27.16	21.24	30.18	33.25	70.63
MLP	21.43	26.46	19.42	27.96	29.83	70.52
Transformer	29.24	30.99	24.26	34.39	35.60	70.57
SCALE _{ft}	29.68	31.83	25.25	36.34	37.38	70.69

tion (and two regression) tasks on minute-long videos. Past studies [31, 60] have established that increasing the input’s time span enhances accuracy in this challenging dataset. As shown in Table 2, our experiments indicate that SCALE can improve the baseline model’s performance by a considerable margin. Moreover, we found that fine-tuning SCALE can lead to further enhancements.

SSv2: Multiple classes in SSv2 share similar backgrounds and only differ in motion [28], suggesting that high performance on this dataset demonstrates that the model has captured strong motion-related contextual cues [44]. Results in Table 3 show that we outperform the state-of-the-art. On this dataset, we see a large performance gap between models that process single clips at a time (Linear and MLP) and the models that work with multiple clips (SCALE and Transformer). We can see SCALE_{linear} is also outperforming the MLP, showing that SCALE is able to capture motion and long-form temporal features of the video. We even improve the supervised model trained on SSv2 ($VMAE_{SSv2}^B$).

UCF-101 & HMDB-51: For these smaller datasets, besides linear and nonlinear probing, we also use k -NN probing (see Table 4 and Table 5). With SCALE _{k -NN}, we see a consistent improvement over the baseline and find that pre-trained MAE-based models greatly benefit from our

Table 4. **UCF Results.** Linear and nonlinear probing accuracies on UCF-101 [48]. SCALE_{ft} outperforms all the other models and, in the case of ρ BYOL, even gets performance close to a fully fine-tuned model. Also, in most cases, SCALE_{linear} outperforms the fine-tuned Transformer and achieves state-of-the-art results in linear probing (previous SotA using RGB frames was 92.6 [45]). We further see a significant accuracy improvement in k -NN probing, especially for pre-trained MAE-based models. As a point of reference, the current best fully fine-tuned accuracy (which is not comparable with our setting) is 96.8% [57].

	SVT	ρ BYOL	VMAE ^B	VMAE ^L	VMAE ^B _{ft}
k -NN	87.20	85.19	35.05	49.14	96.82
SCALE _{k-NN}	89.00	83.47	65.63	76.02	97.38
Linear	91.27	89.55	66.53	84.53	97.91
SCALE _{linear}	<u>92.65</u>	91.43	<u>74.46</u>	86.78	<u>98.14</u>
MLP	91.17	93.60	71.97	<u>87.04</u>	98.04
Transformer	92.20	<u>94.34</u>	68.22	<u>86.30</u>	98.04
SCALE _{ft}	92.94	95.00	76.07	89.92	98.46
FT _{reported}	93.7	95.4	96.1	-	-

Table 5. **HMDB Results.** Linear and nonlinear probing accuracies on HMDB-51 [34]. Despite the small size of the dataset, we see that SCALE_{ft} is outperforming all the other methods, and in the case of SVT, it even outperforms the fully fine-tuned model. We also see that SCALE_{linear} outperforms the Transformer in most cases with only a single linear layer (the best linear accuracy in the literature is 66.7% [45]). Similar to UCF results, we see a considerable increase in the performance of k -NN classifiers for pre-trained MAE-based models.

	SVT	ρ BYOL	VMAE ^B	VMAE ^L	VMAE ^B _{ft}
k -NN	51.83	49.67	21.96	29.21	72.81
SCALE _{k-NN}	56.01	51.56	37.18	51.30	71.83
Linear	63.07	61.17	45.22	60.26	76.33
SCALE _{linear}	<u>66.33</u>	63.92	52.15	62.35	<u>78.36</u>
MLP	63.00	64.77	49.01	<u>62.61</u>	77.45
Transformer	63.98	<u>66.16</u>	47.32	61.50	76.86
SCALE _{ft}	68.10	66.79	<u>51.89</u>	64.83	79.34
FT _{reported}	67.2	73.6	73.3	-	-

training. This can be explained by the additional invariance properties introduced through the SET loss term in SCALE training. Across the board, we also see that in the case of linear probing, not only does SCALE_{linear} outperform Linear, but it also outperforms Transformer, which leverages many more parameters. In the case of SVT, our SCALE_{linear} also outperforms the best reported linear accuracy on UCF101 (92.7% vs. 92.6% [45]). Finally, SCALE_{ft} achieves better results than all the nonlinear baselines and even outperforms the fully fine-tuned SVT (68.1% vs. 67.2%).

Kinetics-400: We present our main results on Kinetics-400 [33] in Table 6. Our SCALE_{linear} with SVT backbone beats the previous state of the art (71.8% vs. 71.5% [18]) and SCALE_{ft} can even improve the accuracy of VMAE^B_{ft},

Table 6. **Kinetics-400 Results.** Linear and nonlinear probing accuracies on Kinetics-400 [33] without any extra data and using RGB frames only. While SCALE_{linear} is on par with Linear, we observe clear improvements for nonlinear probing in the case of SCALE_{ft}. Note that the best linear accuracy on this dataset (without any extra data) is 71.5 [18] and the best full fine-tuning accuracy is 86.7 [61].

	SVT	ρ BYOL	VMAE ^B	VMAE ^L	VMAE ^B _{ft}
Linear	71.71	68.82	43.50	60.73	81.52
SCALE _{linear}	71.78	68.38	43.96	60.66	81.44
MLP	71.19	<u>69.42</u>	<u>45.48</u>	61.64	81.27
Transformer	<u>72.18</u>	69.28	44.85	<u>62.15</u>	<u>81.70</u>
SCALE _{ft}	72.38	69.63	46.15	62.67	81.84

Table 7. **Kinetics-400 Low-shot Results.** Linear and nonlinear probing accuracies on 10% of Kinetics-400 [33]. SCALE is more robust to the size of the labeled dataset. SCALE_{ft} does not overfit like the other nonlinear probes (MLP and Transformer) and outperforms the baselines. VMAE^B_{ft} is greyed out since it was already fine-tuned on the whole labeled dataset and is only reported here for consistency with the other tables.

	SVT	ρ BYOL	VMAE ^B	VMAE ^L	VMAE ^B _{ft}
Linear	<u>66.43</u>	56.43	31.25	48.42	79.79
SCALE _{linear}	65.96	57.74	34.03	<u>49.21</u>	<u>79.94</u>
MLP	65.44	58.68	30.47	48.27	79.37
Transformer	64.97	<u>58.95</u>	29.89	48.27	79.47
SCALE _{ft}	67.01	59.52	<u>33.92</u>	50.36	80.47

which is a strong supervised model, from 81.5% to 81.84%.

Following the evaluation setup of self-supervised image representations [8, 7, 71], we also introduce low-shot K400 video classifications by sampling 10 percent of the videos (in a class-balanced way) and training the probes only on those. We still test on the whole evaluation set of K400. This low-shot setting is more aligned with the typical use-case of self-supervised models in which there is abundant unlabeled data for training via self-supervision and a small set of labeled data for fine-tuning. Results in Table 7 show that our method is particularly effective in this low-shot setting. While most other nonlinear probes overfit and perform worse than the linear probes, our SCALE_{ft} does not overfit and clearly outperforms the baselines.

4.3. Ablations

In this section, we start from a baseline setup consisting of a two-layer transformer with a hidden size of 256, 20% chance of masking clips, trained with a batch size of 512 for 200 epochs, and using two sets of 8 views for representation learning. Using SCALE_{ft}, we explore different loss functions, masking ratios, number of layers, and finally, the number of views during training and testing. All experiments are performed on UCF and HMDB.

Loss Function: As explained in the method section, we have two loss terms, and each of them can be enabled or

Table 8. **Loss Function.** SCALE_{fit} accuracy with different loss function combinations (the masking ratio here is 20%). We can see that having MCM is always beneficial, and the SET loss is almost always helpful. We use both loss terms for our final model.

SET	MCM	UCF-101		HMDB-51	
		SVT	ρ BYOL	SVT	ρ BYOL
✓	✗	91.80	92.20	64.50	63.59
✗	✓	92.01	93.81	62.81	64.05
✓	✓	93.20	92.99	64.57	65.61

Table 9. **Masking Ratio.** SCALE_{fit} accuracy with different masking ratios. We observe best results around 25% similar to NLP models [12] (15%), and different from low-level video models like VideoMAE [54] (90%).

Masking Ratio	UCF-101		HMDB-51	
	SVT	ρ BYOL	SVT	ρ BYOL
0.15	93.18	93.25	64.37	64.83
0.25	93.20	93.81	65.49	65.62
0.35	93.15	93.06	64.18	65.22
0.45	92.96	93.02	63.39	64.96

Table 10. **Transformer Capacity.** SCALE_{fit} accuracy with different model capacities. Having more than one transformer layer and not too few hidden channels is necessary for the best performance.

Hidden Dim	Num Layers	UCF-101		HMDB-51	
		SVT	ρ BYOL	SVT	ρ BYOL
64	1	-	92.62	-	63.16
128	1	-	92.83	-	63.68
256	1	-	92.86	-	64.39
128	2	92.78	93.52	63.26	65.55
256	2	93.20	93.81	65.49	65.62
512	2	92.57	93.66	65.68	64.77
128	3	92.33	93.25	64.83	65.55
256	3	92.75	93.52	65.49	65.16
512	3	92.86	92.93	65.49	64.84

disabled for the pretraining. In Table 8 we show that having both loss terms is better than the individual loss terms.

Masking Ratio: Masking ratio is an important hyperparameter and depends on the data modality, for example, BERT [12] uses 15%, MSN [3] uses 30% (for ViT-Base), MAE [25] uses 75%, and VideoMAE [54] uses 90 to 95% masking. Since our clip representations are somewhat abstract representations of the video, we expect the optimal masking ratio to be close to NLP models rather than video MAEs. We have observed a steady decrease in the pretraining task’s performance with higher masking ratios, so we only tested low masking ratios in Table 9 and found out that 25% is the optimal masking ratio.

Transformer Capacity: We also explore the number of

Table 11. **Number of Views.** SCALE_{fit} accuracy with different numbers of clips and batch sizes. More views lead to consistent improvement, and large batch sizes are not necessary because of the hard negative samples.

Num Views	Batch Size	UCF-101		HMDB-51	
		SVT	ρ BYOL	SVT	ρ BYOL
4 × 2	256	92.65	92.83	64.35	64.24
6 × 2	256	92.70	93.07	64.57	64.37
8 × 2	256	92.80	93.49	64.64	64.63
4 × 2	512	92.67	93.49	64.85	64.50
6 × 2	512	93.18	93.68	64.90	65.35
8 × 2	512	93.20	93.81	65.49	65.62
4 × 2	1024	92.75	93.36	64.77	65.15
6 × 2	1024	92.96	93.57	64.96	65.48
8 × 2	1024	93.07	OOM	65.29	OOM

transformer layers and their hidden size in Table 10. We can see that having more than one transformer layer is necessary for good results and too few hidden channels can hurt performance. However, there is a trade-off, and deeper transformers can lead to worse performance.

Number of Views: Finally, we studied the model performance as we changed the number of views and batch size fed to the model. As can be seen in Table 11, having more views has a large and consistent impact on the performance, and since we have hard negatives for contrastive loss within the video, we are not too reliant on large batch sizes.

5. Conclusion

In this paper, we introduced SCALE, a framework for video representation learning that aggregates the information from multiple clips from the same video. We combine contrastive learning and masked modeling with intuitions from predictive coding to obtain improved global and local representations of clips starting from frozen backbones. We evaluated these features using a wide array of backbones on different action classification and video understanding datasets and achieved strong or state-of-the-art results. The computational efficiency of our method is extremely useful for videos and opens the possibility to a wider group of researchers (with limited computational resources) to work on video representation learning than previously possible. We also believe that working with a set of clips is an interesting direction for representation learning. Finally, as a surprising and maybe alarming observation, even contrastive representations that were trained to be invariant to data augmentations and spatio-temporal crops can be used for contrastive masked modeling. This might be due to benign memorization [2], and understanding why this phenomenon happens might lead to a better understanding of contrastive learning.

6. Acknowledgments

This work was supported by grant 200020_188690 of the Swiss National Science Foundation (SNSF) and an Adobe award.

References

- [1] Hassan Akbari, Linagzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. *ArXiv*, abs/2104.11178, 2021. 1
- [2] Sotiris Anagnostidis, Gregor Bachmann, Lorenzo Noci, and Thomas Hofmann. The curious case of benign memorization. *ArXiv*, abs/2210.14019, 2022. 8
- [3] 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 *European Conference on Computer Vision*, pages 456–473. Springer, 2022. 1, 5, 8
- [4] Shekoofeh Azizi, Basil Mustafa, Fiona Ryan, Zach Beaver, Jana von Freyberg, Jonathan Deaton, Aaron Loh, Alan Karthikesalingam, Simon Kornblith, Ting Chen, Vivek Natarajan, and Mohammad Norouzi. Big self-supervised models advance medical image classification. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3458–3468, 2021. 1
- [5] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *ArXiv*, abs/2106.08254, 2022. 1
- [6] Sagie Benaim, Ariel Ephrat, Oran Lang, Inbar Mosseri, William T Freeman, Michael Rubinstein, Michal Irani, and Tali Dekel. Speednet: Learning the speediness in videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9922–9931, 2020. 2
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. 1, 2, 5, 6, 7
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 5, 7
- [9] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9640–9649, 2021. 2, 5
- [10] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3075–3084, 2019. 1
- [11] Ishan Dave, Rohit Gupta, Mamshad Nayeem Rizve, and Mubarak Shah. Tclr: Temporal contrastive learning for video representation. *Computer Vision and Image Understanding*, 219:103406, 2022. 2
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2, 3, 4, 8
- [13] Xiaoyi Dong, Jianmin Bao, Ting Zhang, Dongdong Chen, Weiming Zhang, Lu Yuan, Dong Chen, Fang Wen, and Nenghai Yu. Bootstrapped masked autoencoders for vision bert pretraining. *ArXiv*, abs/2207.07116, 2022. 2, 3
- [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. *arXiv preprint arXiv:2010.11929*, 2020. 2, 5
- [15] Haoqi Fan, Yanghao Li, Bo Xiong, Wan-Yen Lo, and Christoph Feichtenhofer. Pyslowfast. <https://github.com/facebookresearch/slowfast>, 2020. 5
- [16] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. *arXiv preprint arXiv:2205.09113*, 2022. 1, 2, 3, 4
- [17] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6201–6210, 2019. 5, 6
- [18] Christoph Feichtenhofer, Haoqi Fan, Bo Xiong, Ross Girshick, and Kaiming He. A large-scale study on unsupervised spatiotemporal representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3299–3309, 2021. 1, 2, 3, 5, 6, 7
- [19] Rohit Girdhar, Alaaeldin El-Nouby, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Omnimae: Single model masked pretraining on images and videos. *arXiv preprint arXiv:2206.08356*, 2022. 1, 2, 4
- [20] Rohit Girdhar and Kristen Grauman. Anticipative video transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13505–13515, 2021. 3
- [21] Shreyank N Gowda, Marcus Rohrbach, and Laura Sevilla-Lara. Smart frame selection for action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1451–1459, 2021. 3
- [22] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The” something something” video database for learning and evaluating visual common sense. In *Proceedings of the IEEE international conference on computer vision*, pages 5842–5850, 2017. 2, 5, 6
- [23] 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. *Advances in neural information processing systems*, 33:21271–21284, 2020. 1, 2, 5
- [24] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *Advances in Neural Information Processing Systems*, 33:5679–5690, 2020. 2
- [25] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable

- vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#)
- [26] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. [1](#), [2](#)
- [27] R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. *ArXiv*, abs/1808.06670, 2019. [5](#)
- [28] Kaiqin Hu, Jie Shao, Yuan Liu, Bhiksha Raj, Marios Savvides, and Zhiqiang Shen. Contrast and order representations for video self-supervised learning. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7919–7929, 2021. [6](#)
- [29] Deng Huang, Wenhao Wu, Weiwen Hu, Xu Liu, Dongliang He, Zhihua Wu, Xiangmiao Wu, Mingkui Tan, and Errui Ding. Ascnet: Self-supervised video representation learning with appearance-speed consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8096–8105, 2021. [1](#)
- [30] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. [5](#), [6](#)
- [31] Md. Mohaiminul Islam and Gedas Bertasius. Long movie clip classification with state-space video models. *ArXiv*, abs/2204.01692, 2022. [6](#)
- [32] Simon Jenni and Hailin Jin. Time-equivariant contrastive video representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9970–9980, 2021. [2](#)
- [33] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. [2](#), [5](#), [7](#)
- [34] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *2011 International conference on computer vision*, pages 2556–2563. IEEE, 2011. [2](#), [5](#), [7](#)
- [35] Xingbin Liu, Jinghao Zhou, Tao Kong, Xianming Lin, and Rongrong Ji. Exploring target representations for masked autoencoders. *arXiv preprint arXiv:2209.03917*, 2022. [1](#)
- [36] Yue Liu, Junqi Ma, Yufei Xie, Xuefeng Yang, Xingzhen Tao, Lin Peng, and Wei Gao. Contrastive predictive coding with transformer for video representation learning. *Neurocomputing*, 482:154–162, 2022. [3](#)
- [37] Guillaume Lorré, Jaonary Rabarisoa, Astrid Orcesi, Samia Ainouz, and Stephane Canu. Temporal contrastive pretraining for video action recognition. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 662–670, 2020. [1](#), [2](#), [3](#)
- [38] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European conference on computer vision*, pages 527–544. Springer, 2016. [2](#)
- [39] Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12475–12486, 2021. [2](#)
- [40] Charlie Nash, João Carreira, Jacob Walker, Iain Barr, Andrew Jaegle, Mateusz Malinowski, and Peter Battaglia. Transframer: Arbitrary frame prediction with generative models. *arXiv preprint arXiv:2203.09494*, 2022. [3](#)
- [41] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [1](#), [2](#), [3](#), [5](#)
- [42] Mandela Patrick, Yuki M. Asano, Bernie Huang, Ishan Misra, Florian Metze, João F. Henriques, and Andrea Vedaldi. Space-time crop & attend: Improving cross-modal video representation learning. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10540–10552, 2021. [3](#)
- [43] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. [3](#)
- [44] Kanchana Ranasinghe, Muzammal Naseer, Salman Khan, Fahad Shahbaz Khan, and Michael S Ryoo. Self-supervised video transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2874–2884, 2022. [1](#), [2](#), [3](#), [5](#), [6](#)
- [45] Adria Recasens, Pauline Luc, Jean-Baptiste Alayrac, Luyu Wang, Florian Strub, Corentin Tallec, Mateusz Malinowski, Viorica Pătrăucean, Florent Alché, Michal Valko, et al. Broaden your views for self-supervised video learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1255–1265, 2021. [1](#), [2](#), [3](#), [5](#), [7](#)
- [46] Joshua Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive learning with hard negative samples. *ArXiv*, abs/2010.04592, 2021. [5](#)
- [47] Fadime Sener, Dipika Singhania, and Angela Yao. Temporal aggregate representations for long-range video understanding. In *European Conference on Computer Vision*, pages 154–171. Springer, 2020. [3](#)
- [48] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. [2](#), [5](#), [7](#)
- [49] Chen Sun, Arsha Nagrani, Yonglong Tian, and Cordelia Schmid. Composable augmentation encoding for video representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8834–8844, 2021. [3](#)
- [50] Dídac Surís, Ruoshi Liu, and Carl Vondrick. Learning the predictability of the future. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12607–12617, 2021. [3](#)

- [51] Hao Tan, Jie Lei, Thomas Wolf, and Mohit Bansal. Vimpac: Video pre-training via masked token prediction and contrastive learning. *arXiv preprint arXiv:2106.11250*, 2021. [1](#), [3](#)
- [52] Chenxin Tao, Xizhou Zhu, Gao Huang, Yu Qiao, Xiaogang Wang, and Jifeng Dai. Siamese image modeling for self-supervised vision representation learning. *arXiv preprint arXiv:2206.01204*, 2022. [3](#)
- [53] Nenad Tomasev, Ioana Bica, Brian McWilliams, Lars Holger Buesing, Razvan Pascanu, Charles Blundell, and Jovana Mitrovic. Pushing the limits of self-supervised resnets: Can we outperform supervised learning without labels on imagenet? In *First Workshop on Pre-training: Perspectives, Pitfalls, and Paths Forward at ICML 2022*, 2022. [1](#)
- [54] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [8](#)
- [55] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021. [2](#)
- [56] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [5](#)
- [57] Jue Wang, Gedas Bertasius, Du Tran, and Lorenzo Torresani. Long-short temporal contrastive learning of video transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14010–14020, 2022. [3](#), [7](#)
- [58] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks for action recognition in videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41:2740–2755, 2019. [5](#)
- [59] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Yu-Gang Jiang, Luowei Zhou, and Lu Yuan. Bevt: Bert pretraining of video transformers. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14713–14723, 2022. [1](#), [4](#)
- [60] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. [6](#)
- [61] Chen Wei, Haoqi Fan, Saining Xie, Chaoxia Wu, Alan Loddon Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14648–14658, 2022. [7](#)
- [62] Donglai Wei, Joseph J Lim, Andrew Zisserman, and William T Freeman. Learning and using the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8052–8060, 2018. [2](#)
- [63] Chaoxia Wu and Philipp Krähenbühl. Towards long-form video understanding. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1884–1894, 2021. [2](#), [5](#), [6](#)
- [64] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 284–293, 2019. [3](#)
- [65] Chao-Yuan Wu, Yanghao Li, Karttikeya Mangalam, Haoqi Fan, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Memvit: Memory-augmented multiscale vision transformer for efficient long-term video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13587–13597, 2022. [1](#), [3](#)
- [66] Yu Wu, Linchao Zhu, Xiaohan Wang, Yi Yang, and Fei Wu. Learning to anticipate egocentric actions by imagination. *IEEE Transactions on Image Processing*, 30:1143–1152, 2020. [3](#)
- [67] Shen Yan, Xuehan Xiong, Anurag Arnab, Zhichao Lu, Mi Zhang, Chen Sun, and Cordelia Schmid. Multiview transformers for video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3333–3343, 2022. [3](#)
- [68] Liangzhe Yuan, Rui Qian, Yin Cui, Boqing Gong, Florian Schroff, Ming-Hsuan Yang, Hartwig Adam, and Ting Liu. Contextualized spatio-temporal contrastive learning with self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13977–13986, 2022. [3](#)
- [69] Yujia Zhang, Lai-Man Po, Xuyuan Xu, Mengyang Liu, Yexin Wang, Weifeng Ou, Yuzhi Zhao, and Wing-Yin Yu. Contrastive spatio-temporal pretext learning for self-supervised video representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 3380–3389, 2022. [3](#)
- [70] Zhang Zhang and Dacheng Tao. Slow feature analysis for human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 34(3):436–450, 2012. [1](#)
- [71] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021. [1](#), [2](#), [3](#), [7](#)
- [72] Pan Zhou, Yichen Zhou, Chenyang Si, Weihao Yu, Teck Khim Ng, and Shuicheng Yan. Mugs: A multi-granular self-supervised learning framework. *arXiv preprint arXiv:2203.14415*, 2022. [1](#)