

# How Severe is Benchmark-Sensitivity in Video Self-Supervised Learning?

Fida Mohammad Thoker, Hazel Doughty, Piyush Bagad, Cees Snoek

University of Amsterdam

**Abstract.** Despite the recent success of video self-supervised learning models, there is much still to be understood about their generalization capability. In this paper, we investigate how sensitive video self-supervised learning is to the current conventional benchmark and whether methods generalize beyond the canonical evaluation setting. We do this across four different factors of sensitivity: domain, samples, actions and task. Our study which encompasses over 500 experiments on 7 video datasets, 9 self-supervised methods and 6 video understanding tasks, reveals that current benchmarks in video self-supervised learning are not good indicators of generalization along these sensitivity factors. Further, we find that self-supervised methods considerably lag behind vanilla supervised pre-training, especially when domain shift is large and the amount of available downstream samples are low. From our analysis we distill the *SEVERE-benchmark*, a subset of our experiments, and discuss its implication for evaluating the generalizability of representations obtained by existing and future self-supervised video learning methods. Code is available at <https://github.com/fmthoker/SEVERE-BENCHMARK>.

**Keywords:** Self-supervised learning, Video representation learning, Video understanding

## 1 Introduction

Video self-supervised learning has progressed at a tremendous pace in recent years, *e.g.* [1, 54, 56–58, 75], as it offers a crucial starting point from which to learn. This is especially important for video understanding applications, where annotating large amounts of data is extremely expensive, error-prone and sensitive to annotator bias. Hence, learning video representations through self-supervision is crucial, especially for use cases where the downstream video data is limited because of the domain, task or actions the video contains. However, the majority of current works in video self-supervised learning, *e.g.* [4, 48, 49, 53, 81], do not test beyond standard benchmarks. The standard protocol is to use unlabeled Kinetics-400 [36] for pre-training and then measure performance by fine-tuning on two action recognition datasets: UCF-101 [65] and HMDB-51 [42]. While these benchmarks have facilitated the impressive progress of video self-supervised learning in recent years, they cannot indicate the generalizability of such methods as these pre-training and downstream datasets are all similar in appearance and the type of actions they contain. Some methods have

started to report finetuning performance on additional datasets like Something-Something-v2 [25] in [20, 56, 75], Diving-48 [43] in [14, 78], AVA [27] in [20, 80, 82] and EPIC-Kitchens-100 [13] in [82]. However, such evaluations are insufficient to understand the generalization of video self-supervised methods alone since they only add a single additional dataset, often without comparison to prior methods.

In this work, we address the essential need to gauge the sensitivity of existing video self-supervised methods to the current benchmark by thoroughly evaluating their performance for generalization across diverse downstream settings. Similar benchmarking studies have been performed for self-supervised pre-training in images [5, 12, 16, 17, 24, 33, 38, 41, 50, 60, 73, 83, 86], which investigate model transferability [16, 33, 50, 74] or the importance of factors like pre-training dataset [12, 24, 41] and backbone architecture [38]. Unfortunately, lessons from these works do not directly transfer to video self-supervised learning. First, video self-supervised tasks are distinct from those of images as they are designed to understand the temporal dimension of video [14, 56, 75, 82] in addition to the spatial understanding needed in images [9]. Second, video is multi-modal and several methods [4, 49, 54] are designed to exploit cross or multi-modal understanding, which is again absent in image-based methods. For videos, [20] extends four image-based self-supervised methods to videos and investigate their performance focusing on different pre-training setups. We take inspiration from this and benchmarking works in image self-supervised learning and perform a much-needed study for understanding the generalizability of self-supervised methods for video in relation to different downstream factors.

As our first contribution, we identify the problem of benchmark-sensitivity in video self-supervised learning and examine this sensitivity along the factors of domain, samples, actions and task. As our second contribution, we perform an extensive evaluation which spans a total of over 500 experiments with 9 video self-supervised learning methods across 7 video datasets and 6 video understanding tasks. We find that standard benchmarks in video self-supervised learning do not indicate generalization along the said sensitivity factors and vanilla supervised pre-training outperforms self-supervised pre-training, particularly when domain change is large and there are only a few downstream finetuning samples available. Third, we propose a subset of our experiments as the SEVERE-benchmark for future self-supervised learning methods to benchmark generalization capability. We also discuss the implication of this benchmark for evaluating the generalizability of representations obtained by existing methods as well as the nature of video self-supervised objectives that currently generalize well.

## 2 Identifying Benchmark Sensitivity

The vast majority of current works in video self-supervised learning evaluate their approach by pre-training on Kinetics-400 [36] and finetuning the learned representation for action recognition on UCF-101 [65] and HMDB-51 [42]. Some works [4, 14, 22, 31, 44, 54, 56, 70, 75] also report performance on video retrieval for UCF-101 and HMDB-51 and several recent works [58, 59, 82] compare linear

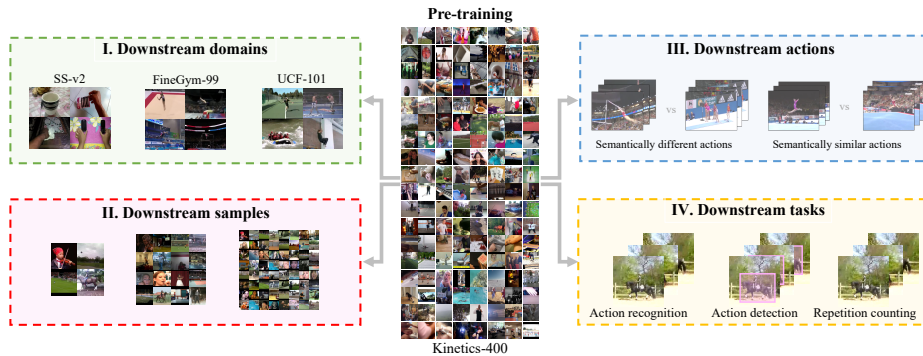


Fig. 1: **Benchmark-sensitivity.** We evaluate the sensitivity of 9 video self-supervised learning methods along 4 downstream factors which vary from the pre-training source: the domain, the samples, the actions and the task.

evaluation performance on Kinetics-400. However, these downstream datasets are very similar to each other and also share many similarities with the pre-training dataset of Kinetics-400. Videos in all three datasets are collected from YouTube and are mostly recorded with a single camera containing a single well-positioned human actor. In terms of class labels, all datasets focus on similar, coarse-grained and mutually exclusive actions with many actions common between pre-training and downstream datasets. Besides all these data similarities, the existing evaluations also ignore a major benefit of self-supervised representation learning for videos, *i.e.* finetuning the representation with only a small amount of data samples and transferring to other video understanding tasks beyond action recognition. Hence, we believe the current benchmark standard is insufficiently equipped to gain a true understanding of where video self-supervised models are successful, as it cannot show the generalizability or the sensitivity of methods to factors such as domain shift, amount of finetuning data samples, action similarity or task shift. In this study, we identify the sensitivity of existing evaluations and thoroughly benchmark self-supervised video learning methods along four sensitivity factors as depicted in Fig. 1.

- I. **Downstream domain.** First, we analyse whether features learned by self-supervised models transfer to datasets that vary in domain with respect to the pre-training dataset.
- II. **Downstream samples.** Second, we evaluate the sensitivity of self-supervised methods to the number of downstream samples available for finetuning.
- III. **Downstream actions.** Third, we investigate if self-supervised methods learn fine-grained features required to recognize semantically similar actions.
- IV. **Downstream task.** Finally, we study the sensitivity of video self-supervised methods to the downstream task and question whether self-supervised features can be used beyond action recognition.

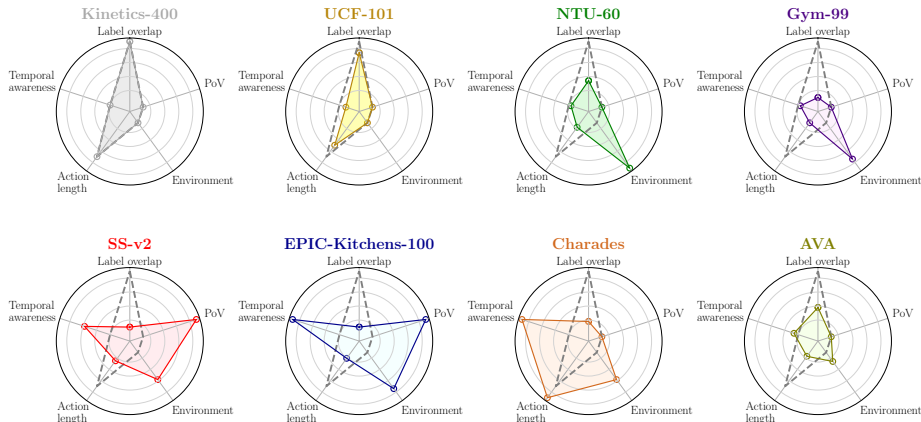


Fig. 2: **Video dataset characteristics.** Characterizing domain shift in datasets via difference in label overlap, point-of-view (PoV), environment, action length and temporal awareness with Kinetics-400 (shown by dotted line). Kinetics-400 and UCF-101 are highly similar to each other, while datasets like Something-Something-v2, EPIC-Kitchens-100 and Charades have different attributes compared to Kinetics-400.

## 2.1 Downstream Video Datasets

We evaluate various self-supervised models along our four sensitivity factors on 7 video datasets: **UCF-101** [65], **NTU-60** [62], **FineGym** (Gym-99) [63], **SomethingSomething-v2** (SS-v2) [25], **EPIC-Kitchens-100** (EK-100) [13], **Charades** [64] and **AVA** [27]. They include a considerable variety in video domain, the actions they contain and cover a range of video understanding tasks. To get a sense of the differences between these downstream datasets and the Kinetics-400 source dataset, we summarize their similarity to Kinetics-400 by radar plots in Fig. 2 based on several attributes. *Environment* refers to the variety of settings contained in the dataset. *Point-of-view* is whether a video is recorded from a first-person or third-person viewpoint. *Temporal awareness* defines the extent to which temporal context is required to recognize or detect actions. We quantify this as the point at which performance saturates with increasing temporal context in the input. *Label overlap* is the fraction of actions in a target dataset that are also present in Kinetics-400. *Action length* is the temporal length of the actions in seconds. Details are provided in the appendix.

## 2.2 Evaluated Self-Supervised Video Learning Methods

Self-supervised learning methods in video can be grouped into two categories based on the objective they use: pretext task methods and contrastive learning methods. Pretext task methods use predictive tasks such as solving spatio-temporal jigsaw puzzles [2, 32, 37], rotation prediction [35], frame and clip order [21, 48, 68, 81, 84], video speed [7, 11, 34, 77, 85], video completion [45], predicting motion statistics [76], tracking random patches in video frames [75] or

audio-visual clustering [3, 4, 8, 30]. Contrastive learning methods discriminate between ‘positive’ and ‘negative’ pairs to learn invariances to certain data augmentations and instances either from visual-only input [14, 15, 28, 44, 53, 58, 66, 82] or multi-modal data [29, 40, 46, 49, 54, 69, 71].

Some methods also combine pretext and contrastive approaches [6, 15, 31, 56, 70, 88]. A detailed survey of video self-supervised learning methods can be found in [61]. We consider 9 video-based self-supervised methods which achieve good performance on current benchmarks and cover a range of self-supervised paradigms in the video domain, including contrastive learning, pretext-tasks, their combination and cross-modal audio-video learning.

Due to the high computational cost of training self-supervised methods, we focus on works with publicly available weights for a common R(2+1)D-18 network [72] pre-trained on Kinetics-400 [36]: **MoCo** [10], **SeLaVi** [4], **Video-MoCo** [53], **Pretext-Contrast** [70], **RSPNet** [56], **AVID-CMA** [49], **CtP** [75], **TCLR** [14] and **GDT** [54]. We compare these to no pre-training, *i.e.* training from scratch, and fully supervised pre-training for action recognition. It is worth noting that since we use publicly available models we cannot control the exact pre-training setup. There are subtle differences in the training regime for each method, such as the number of epochs, the data augmentations used and the batch size. Details of these differences are provided in the appendix. However, all models use the same backbone and pre-training dataset thus we can evaluate their downstream abilities in exactly the same way. To finetune for downstream tasks we simply attach a task-dependent head at the last layer of the pre-trained R(2+1)D-18 backbone to produce label predictions for the corresponding task. For a fair comparison, we use the same set of hyper-parameters, optimization and pre-processing during the downstream training of each model.

### 3 Sensitivity Factor I: Downstream Domain

We first investigate to what extent self-supervised methods learn features that are applicable to action recognition in any domain. We evaluate the suite of pre-trained models on UCF-101, NTU-60, Gym-99, SS-v2 and EK-100 for the task of action recognition. It is worth noting that as well as variety in domain, these datasets include variety in the amount of training data (9.5k - 168k examples) and cardinality of classification (60 - 300 classes). We attach a single classification layer to the pre-trained backbone and evaluate the models’ performance on the downstream task in two settings. First, **full finetuning** where we train the whole network from the initialization of the pre-trained weights. Second, **linear evaluation** where we train the classification layer only using the frozen features of pre-trained backbones. We follow the standard splits proposed in the original datasets and report video-level top-1 accuracy on the test sets. The details about splits, pre-processing, training for each dataset are provided in the appendix.

**Full finetuning.** The left part of Table 1 shows the results of full finetuning. From the results, it is clear that all self-supervised methods are very effective on UCF-101 as there is a significant gap between training from scratch and

Table 1: **Sensitivity Factor I: Downstream Domain.** Video self-supervised methods evaluated across datasets with increasing domain shift with respect to the source dataset (see Fig. 2). Colors denote relative rankings across methods for each dataset, ranging from low (red) to high (blue). The ranking of methods is domain-sensitive for both finetuning and linear classification and becomes less and less correlated with the current UCF-101 benchmark as the domain shift increases.

Pre-training	Finetuning					Linear Evaluation					
	UCF101	NTU60	Gym99	SSv2	EK 100	K 400	UCF101	NTU60	Gym99	SSv2	EK 100
None	77.3	92.9	89.8	57.1	25.7	-	-	-	-	-	-
MoCo	83.3	93.4	90.7	57.1	26.4	34.5	65.4	16.0	21.2	7.4	21.4
VideoMoCo	84.9	94.1	90.3	59.0	43.6	31.0	66.3	51.6	41.6	19.5	25.7
SeLaVi	85.2	92.8	88.9	56.2	33.8	24.1	51.2	15.7	20.2	4.5	22.4
Pretext-Contrast	87.7	93.9	90.5	56.9	34.3	22.4	57.2	17.6	30.0	10.9	20.0
RSPNet	88.7	93.9	91.1	59.0	42.7	46.0	76.6	33.5	32.2	12.5	24.9
AVID-CMA	88.8	94.0	90.4	52.0	29.9	43.5	78.1	53.9	45.1	16.1	22.5
CtP	90.1	94.3	92.0	59.6	42.8	7.6	37.9	22.6	30.6	12.2	20.0
TCLR	90.8	94.1	91.6	59.8	36.2	19.9	63.3	33.5	33.0	10.8	21.8
GDT	91.3	93.9	90.5	58.0	37.3	38.6	75.7	38.2	34.2	11.9	25.3
Supervised	93.9	93.9	92.1	60.8	47.7	65.9	91.7	45.5	42.7	16.6	26.6

using self-supervised pre-training. This gap is reduced as the difference between Kinetics-400 and the downstream domain increases. SeLaVi, MoCo and AVID-CMA in particular are evidence of this as these methods suffer when datasets have higher temporal awareness and less label overlap with Kinetics-400. When moving from UCF-101 to NTU-60 and Gym-99 there is a change in the ordering of self-supervised methods. This demonstrates a high performance on UCF-101 does not guarantee a self-supervised model is generalizable to other domains. The change in ranking is even more prominent for SS-v2 and EK-100, which require the most temporal awareness and also shift to a first-person viewpoint. This is particularly noticeable for AVID-CMA. On these datasets, MoCo has similar results to no pre-training, which is evidence that video-specific self-supervised learning methods are needed and that image-based methods are insufficient. Overall, supervised pre-training achieves good performance across the board, outperforming self-supervised methods on the most similar domain (UCF-101) as well as the most dissimilar domains (SS-v2 and EK-100). Amidst the models tested, CtP, RSPNet, VideoMoCo and TCLR stand out as the self-supervised pre-training methods most generalizable to different domains.

**Linear classification.** The right part of Table 1 shows the results for linear classification. As with finetuning, the ranking among the self-supervised methods changes as the domain difference between the pre-training and the downstream dataset increases. For example, VideoMoCo ranks lower than GDT and RSPNet for UCF-101 and Kinetics-400 but ranks higher than both for all other datasets. This again demonstrates that performance on UCF-101 does not give a complete picture of a self-supervised model’s success. We also observe that linear evaluation on Kinetics-400, as some papers report [58, 59, 82], has the same issue since it is highly correlated to UCF-101 performance. For UCF-101 and Kinetics-400, self-supervised models with contrastive objectives learn highly discriminative

features compared to the non-contrastive models. This can be seen by comparing contrastive models AVID-CMA, GDT and RSPNet to non-contrastive SeLaVi and CtP. From the NTU-60 and Gym-99 results we observe that as the label overlap between the pre-training and the downstream dataset decreases, the performance gap between finetuning and linear evaluation increases considerably. This is true for both supervised and self-supervised pre-training. The most generalizable methods in the linear classification setting are contrastive methods VideoMoCo and AVID-CMA as well as supervised pre-training. Interestingly, there are cases where VideoMoCo and AVID-CMA even outperform supervised pre-training, namely for NTU-60, Gym-99 and SS-v2.

*Conclusion.* We observe from Table 1 that performance for both UCF-101 finetuning and Kinetics-400 linear evaluation is not indicative of how well a self-supervised video model generalizes to different downstream domains, with the ranking of methods changing substantially across datasets and whether full finetuning or linear classification is used.

## 4 Sensitivity Factor II: Downstream Samples

The previous section analyzed sensitivity to the downstream domain by evaluating performance on several different datasets. However, finetuning on each of these datasets uses a large number of labeled examples, which means training from scratch already obtains good performance. Not all domains and use cases have ample labeled video examples available, thus we investigate what the impact of the number of finetuning samples is and whether self-supervised methods can be beneficial in scenarios where we have little data to finetune with. We vary the amount of finetuning data, beginning from 1000 videos, sampled uniformly from the classes, and double the amount until we reach the full training set size. We report on four of the downstream datasets from the previous section: UCF-101, NTU-60, Gym-99 and SS-v2. The results are summarized in Fig. 3.

We first observe that the trends in the low data regime are different from those with the full data. The gap between supervised and self-supervised pre-training is much larger in low data settings, particularly for UCF-101 and Gym-99. NTU is an exception, where, with 1000-4000 samples CtP, GDT, AVID-CMA and TCLR outperform supervised pre-training. As with changes in the downstream domain, change in the amount of downstream examples also causes a change in the ranking of self-supervised models. For example, on UCF-101, RSPNet is much more successful than CtP and TCLR when using only 1000 samples. This is because some self-supervised models benefit more than others from an increased amount of downstream samples. For example, CtP is one of the most generalizable pre-training strategies when finetuning with the full data on UCF-101, Gym-99 and SS-v2, but this is not the case with fewer training samples. Interestingly, GDT is consistently high in the ranking with low amounts of finetuning samples. This is likely due to the large number of temporal augmentations it uses, which help the generalization ability when the training data is limited.

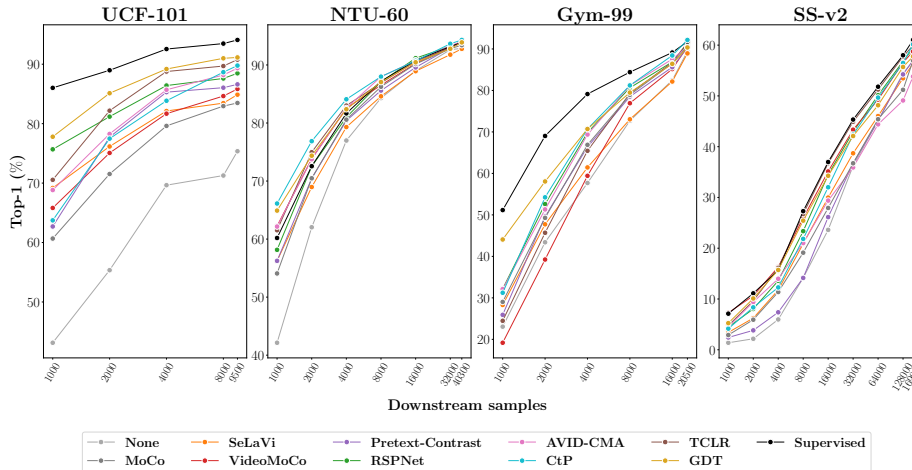


Fig. 3: **Sensitivity Factor II: Downstream Samples.** Comparison of video self-supervised learning methods using varying number of finetuning samples for four downstream datasets. Both the gap and rank among pre-training methods are sensitive to the number of samples available for finetuning.

*Conclusion.* We observe from Fig. 3 that video self-supervised models are highly sensitive to the amount of samples available for finetuning, with both the gap and rank between methods changing considerably across sample sizes on each dataset.

## 5 Sensitivity Factor III: Downstream Actions

As indicated earlier, existing evaluations of self-supervised video learning methods have been limited to coarse-grained action recognition. In this section, we investigate whether current self-supervised tasks are only effective for these types of benchmarks or whether they are able to learn features that are useful for differentiating more challenging and semantically similar actions.

FineGym [63] provides us with an experimental setup to study sensitivity to this factor. The dataset contains different evaluations with varying levels of semantic similarity, namely action recognition *across all events*, *within an event* or *within a set*. Recognition *across all events* uses the whole of Gym-99 containing actions from four gymnastic events. For recognition *within an event* there are two subsets: Vault and Floor containing only actions from these two events. Recognition *within a set* has two subsets namely FX-S1, containing different *leaps-jumps-hops* in Floor, and UB-S1, which consists of types of *circles* in Uneven Bars. We also experiment with the long-tailed version of FineGym, Gym-288, which adds 189 more tail classes. Details of these subsets are in the appendix. As before, we attach a classification head to the pre-trained models



Table 2: **Sensitivity Factor III: Downstream Actions.** Video self-supervised models evaluated on different semantic similarities of action in FineGym: across events, within an event and within a set. Colors denote relative rankings across methods for each dataset, ranging from low (red) to high (blue). Many methods struggle on the within a set benchmark where actions are most semantically similar.

Pre-training	Gym99					Gym288
	Across Events	Within Event		Within Set		Across Events
	All	Vault	Floor	FX-S1	UB-S1	All
None	84.8	24.7	75.9	46.6	82.3	50.0
SeLaVi	84.5	25.4	76.0	51.3	80.9	52.8
AVID-CMA	85.7	30.4	82.7	68.0	87.3	52.5
VideoMoCo	85.9	28.4	79.5	57.3	83.9	54.1
Pretext-contrast	86.0	28.5	81.4	66.1	86.1	52.7
MoCo	86.5	33.2	83.3	65.0	84.5	55.1
GDT	86.6	36.9	83.6	66.0	83.4	55.4
RSPNet	86.9	33.4	82.7	65.4	83.6	55.2
TCLR	87.7	29.8	84.3	60.7	84.7	55.4
CtP	88.1	26.8	86.2	79.1	88.8	56.5
Supervised	88.6	37.7	86.1	79.0	87.1	58.4

and finetune the whole network with the training set of each subset. In Table 2 we report Top-1 accuracy (mean per-class) on the testing sets following [63].

Performance of self-supervised methods varies considerably across downstream actions. The methods that perform best on Gym-99 often do not generalize well to the subsets with higher semantic similarity among actions. This is particularly noticeable for RSPNet and TCLR which drop in the ranking for the within-set subsets. All self-supervised methods, except GDT, struggle on Vault, likely due to the intense motions. Surprisingly, MoCo performs reasonably well when actions are more semantically similar, and is comparable to GDT and RSPNet. The best self-supervised method for subsets with high semantic similarity is CtP. This is especially evident from FX-S1 where it outperforms the second-best self-supervised method, AVID-CMA, by 12%. As with downstream domain and samples, supervised pre-training generalizes better than self-supervised methods across downstream actions with only CtP achieving comparable performance.

Table 2 also compares balanced Gym-99 with long-tailed Gym-288. We observe that self-supervised methods are not robust to this change in distribution, with the gap in performance with respect to supervised pre-training increasing. However, the ranking remains consistent, meaning the performance on the balanced set is generally indicative of the performance on the long-tailed set.

*Conclusion.* Most self-supervised methods in Table 2 are sensitive to the actions present in the downstream dataset and do not generalize well to more semantically similar actions. This further emphasizes the need for proper evaluation of self-supervised methods beyond current coarse-grained action classification.

Table 3: **Sensitivity Factor IV: Downstream Tasks.** Transferability of self-supervised video learning methods across video understanding tasks. Colors denote relative rankings across methods for each task, ranging from low (red) to high (blue). Note that for repetition counting lower (error) is better. Self-supervised features are transferable to different downstream tasks when the domain shift is low, but struggle when there is also a domain shift. Action recognition on UCF-101 is not a good proxy for self-supervised video learning use cases where a downstream domain- and task-shift can be expected.

Pre-training	Task-shift within domain				Task-shift out of domain	
	Action Recognition	Action Detection	Repetition Counting	Arrow of Time	Multi-label Recognition	Action Detection
None	77.3	0.327	0.217	56.1	7.9	7.4
MoCo	83.3	0.416	0.208	80.3	8.3	11.7
VideoMoCo	84.9	0.440	0.185	72.9	10.5	13.1
SeLaVi	85.2	0.419	0.162	77.4	8.4	10.2
Pretext-contrast	87.7	0.462	0.164	77.2	8.9	12.7
RSPNet	88.7	0.467	0.145	87.0	9.0	14.1
AVID-CMA	88.8	0.435	0.148	83.3	8.2	10.0
CtP	90.1	0.465	0.178	77.1	9.6	10.0
TCLR	90.8	0.476	0.142	85.6	12.2	10.8
GDT	91.3	0.463	0.123	76.4	8.5	12.6
Supervised	93.9	0.482	0.132	77.0	23.5	17.9

## 6 Sensitivity Factor IV: Downstream Tasks

The fourth factor we investigate is whether self-supervised video models are sensitive to the downstream task or whether features learned by self-supervised models are useful to video understanding tasks beyond action recognition. We evaluate this in two ways. First, we keep the domain fixed and evaluate different tasks in a domain similar to the pre-training dataset. We also explore further tasks by changing the domain and seeing how these two factors interplay.

**Task-shift within domain.** We consider three different tasks which are all defined for UCF-101: spatio-temporal action detection [39], repetition counting [87] and arrow-of-time prediction [23]. Using UCF-101 allows us to keep the domain fixed across tasks and eliminates the impact of domain shift. Note that each task uses a different subset of the full UCF-101 dataset, however, the domain remains consistent. For each task, we use the R(2+1)D-18 networks as the pre-trained backbones as before and attach task-dependent heads. We report mean Average Precision for spatio-temporal localization [47], mean absolute counting error for repetition counting [87] and classification accuracy for arrow-of-time prediction [23, 79]. Further details are in the appendix.

From the results in Table 3, we observe that self-supervised learning is beneficial to tasks beyond action recognition, with almost all methods outperforming training from scratch on spatio-temporal action detection, repetition counting and arrow-of-time prediction. Action detection results are well correlated with action recognition. Repetition counting and arrow-of-time have less correlation

with action recognition, suggesting that the current benchmark on UCF-101 action recognition by itself is not a good indication of how well self-supervised methods generalize to other tasks. For repetition counting and arrow-of-time prediction, some methods perform comparably to or outperform supervised pre-training. Notably, RSPNet and TCLR generalize the best across these tasks, with GDT also performing well on repetition counting. CtP ranks high on action recognition and detection but performs modestly for repetition counting. This shows that different methods have different task sensitivity, so a thorough evaluation along downstream tasks is needed.

**Task-shift out of domain.** We also evaluate how well the self-supervised models generalize when both the domain and the task change. We do so with two popular video understanding benchmarks: long-term multi-label classification on Charades [64] and short-term spatio-temporal action detection on AVA [27]. For both, we follow the setup and training procedure from [19] with R(2+1)D-18 models as the pre-trained backbone and we measure performance in mean Average Precision. Details are in the appendix.

From the results in Table 3, we observe that supervised pre-training is far more generalizable than all self supervised methods, which all struggle considerably when both the domain and task change. For long-term action classification on Charades, TCLR is slightly better than other methods. On AVA, RSPNet is the best performing self-supervised method with VideoMoCo second. In Section 3, we earlier observed that these were two of the methods more robust to domain shift suggesting that this factor is key to success on AVA.

*Conclusion.* The results in Table 3 reveal that action classification performance on UCF-101 is mildly indicative for transferability of self-supervised features to other tasks on UCF-101. However, when methods pre-trained on Kinetics-400 are confronted with a domain change in addition to the task change, UCF-101 results are no longer a good proxy and the gap between supervised and self-supervised pre-training is large.

## 7 SEVERE-benchmark

As evident from the results in previous sections, current video self-supervised methods are benchmark-sensitive to the four factors we have studied. Based on our findings, we propose the SEVERE-benchmark (Sensitivity of Video Representations) for use in future works to more thoroughly evaluate new video self-supervised methods for generalization along the four sensitivity factors we have examined. Since we do not expect future works to run all the experiments from our study, we create a subset of experiments that are indicative benchmarks for each sensitivity factor and realistic to run. We summarize the benchmark composition in Table 4 and detail its motivation per factor. Standard deviations for the results we obtain on this benchmark can be found in the appendix.

Table 4: **Proposed SEVERE-benchmark** for evaluating video self-supervised methods for generalization along downstream domains, samples, actions and tasks.

Pre-training	Existing		SEVERE-benchmark						
	UCF101	Domains		Samples		Actions		Tasks	
		SS-v2	Gym-99	UCF (10 <sup>3</sup> )	Gym-99 (10 <sup>3</sup> )	FX-S1	UB-S1	UCF-RC	Charades-MLC
None	77.3	57.1	89.8	38.3	22.7	46.6	82.3	0.217	7.9
MoCo	83.3	57.1	90.7	60.4	30.9	65.0	84.5	0.208	8.3
VideoMoCo	84.9	59.0	90.3	65.4	20.6	57.3	83.9	0.185	10.5
SeLaVi	85.2	56.2	88.9	69.0	30.2	51.3	80.9	0.162	8.4
Pretext-Contrast	87.7	56.9	90.5	64.6	27.5	66.1	86.1	0.164	8.9
RSPNet	88.7	59.0	91.1	74.7	32.2	65.4	83.6	0.145	9.0
AVID-CMA	88.8	52.0	90.4	68.2	33.4	68.0	87.3	0.148	8.2
CtP	90.1	59.6	92.0	61.0	32.9	79.1	88.8	0.178	9.6
TCLR	90.8	59.8	91.6	72.6	26.3	60.7	84.7	0.142	12.2
GDT	91.3	58.0	90.5	78.4	45.6	66.0	83.4	0.123	8.5
Supervised	93.9	60.8	92.1	86.6	51.3	79.0	87.1	0.132	23.5

**Downstream domain.** To measure a self-supervised model’s domain sensitivity we recommend using Something-Something-v2 and FineGym-99. These two datasets come from domains distinct to Kinetics-400 and UCF-101 and also each other. FineGym-99 evaluates a model’s ability to generalize to datasets with less distinctive backgrounds where there are few actions in common with Kinetics-400. SS-v2 evaluates the generalizability to actions that require high temporal awareness as well as the shift to a first-person viewpoint. It is evident from Table 4 that there are significant rank changes between UCF-101, Gym-99 and SS-v2 thus these three datasets provide a challenging subset for future methods.

**Downstream samples.** For the sample sensitivity, we recommend using 1000 samples on UCF-101 and Gym-99. Using 1000 samples showed the most dramatic difference from the full dataset size particularly for these datasets where there is a considerable gap between self-supervised and supervised pre-training as well as considerable rank change among the methods.

**Downstream actions.** To test generalizability to recognizing semantically similar actions, we recommend evaluating the two within-set granularities of Gym-99 *i.e.* FX-S1 and UB-S1. Both of these subsets have high semantic similarity between actions with methods currently struggling to generalize to both of these subsets as can be seen in Table 4. There is also a significant gap between supervised and most self-supervised pre-training methods for FX-S1, highlighting the potential for future works in this area.

**Downstream task.** To evaluate the task sensitivity, we recommend using repetition counting on UCF-101 and multi-label classification on Charades. Repetition counting on UCF-101 highlights different strengths to action recognition as it allows investigation of a model’s ability to generalize to a task that requires more temporal understanding without measuring the impact of the domain. We recommend multi-label classification on Charades as it is currently a very challenging task for self-supervised models and allows the combination of domain and task shift to be investigated. Code to compare on the SEVERE-benchmark is available at <https://github.com/fmthoker/SEVERE-BENCHMARK>.

## 8 Observations, Limitations and Recommendations

**Observations.** We hope that our study and resulting benchmark provides a helpful insight for future research to design novel self-supervised methods for generalizable video representation learning. From the benchmark results in Table 4, we observe that:

- (i) There is no clear winner as different methods stand out in different downstream settings.
- (ii) Supervised pre-training is dominant across all sensitivity factors, especially when the number of available downstream samples are limited and when there is a change in both the downstream domain and the downstream task.
- (iii) Self-supervised contrastive methods that explicitly encourage features to be distinct across the temporal dimension transfer well. This is visible from the consistent performance of GDT, TCLR and RSPNet across different sensitivity factors.
- (iv) Learning certain temporal invariances may prevent generalizability to temporal or fine-grained benchmarks. This is evident from GDT’s performance on SS-v2 and UB-S1. These benchmarks require distinction between actions such as *moving something left* vs. *moving something right* in SS-v2 and *giant circle forwards* vs. *giant circle backwards* in UB-S1. The invariance to temporal reversal learned by GDT impacts its ability to recognize such actions. Similarly, MoCo outperforming VideoMoCo on the FX-S1 and UB-S1 Gym-99 subsets suggests that invariance to frame dropout in VideoMoCo can harm the performance on highly similar actions.
- (v) Pretext-tasks specific to videos can be effective to learn more fine-grained features. CtP generalizes well both to different domains where the background is less indicative of the action and to more semantically similar actions. The pretext task is to track and estimate the position and size of image patches moving in a sequence of video frames. Such a formulation requires the network to learn to follow moving targets and ignore the static background information. CtP’s generalization success demonstrates that contrastive learning is not the only way forward for self-supervised video representation learning.
- (vi) Fig. 4 shows the feature similarity on Kinetics using centered kernel alignment [52] between supervised pre-training and the best self-supervised methods *i.e.* GDT, RSPNet, TCLR, CtP. This figure illustrates that contrastive methods seem to imitate supervised pre-training as the correlation between supervised pre-training and the three contrastive methods (RSPNet, GDT and TCLR) is high. This explains the good performance of these methods on UCF-101 with 1000 examples. By contrast, CtP’s features are far away from supervised pre-training. This is interesting because CtP generalizes well to new domains and actions, it shows that good generalization capability can be obtained without imitating supervised pre-training.

**Limitations.** While our study has highlighted the benchmark sensitivity of video self-supervised learning across four factors, there are many more factors




Fig. 4: **Representation similarity** between features of top self-supervised methods and supervised pre-training on Kinetics-400 validation set (using centered kernel alignment [52]). Contrastive methods have a high correlation with supervised pretraining, while CtP’s features are far away. Thus, showing potential for both imitating supervised learning as well as learning features distinct to it.

that we do not consider in this work. Due to computational limits, we keep the source dataset fixed as Kinetics-400 and use publicly available pre-trained models. This means there is variability in the exact pre-training setup such as the spatial data augmentations that are used by each model. We hope that future works will explore impact of such pretraining factors as well as the impact of pre-training on other large-scale datasets such as Ego4D [26] for the generalization of video self-supervised models. Another limitation of our study is that we only consider a fixed R(2+1)D-18 backbone, which is currently one of the most commonly used in video self-supervised learning. This allows our comparison between methods to be fair, however, it does limit the ability of methods to perform well on datasets such as EPIC-Kitchens-100. Another factor that could be explored further is the task. We have considered a selection of various video understanding tasks centered around human actions. However, there are many more video understanding tasks that could be explored such as human centric tasks like action anticipation [13] and temporal action detection [13], as well as non-human centric tasks like animal behavior analysis [18, 51, 67], multi-object tracking [55] and visual grounding [67].

**Recommendations.** Based on the results and our observations, we have several recommendations for future works in video self-supervised learning. (i) Our study has highlighted the need for more focus on generalizability of self-supervised learning methods, particularly along the domain and dataset size factors. (ii) Distinguishing across the temporal dimension is effective and is a useful direction to pursue further for generalizability. (iii) Pretext-tasks like the one used in CtP are good for the generalizability to domain and action, thus designing new video specific pretext tasks is a promising direction. This could also be combined with contrastive learning tasks to gain the benefits of both types of learning.

**Acknowledgements.** This work is part of the research programme Perspectief EDL with project number P16-25 project 3, which is financed by the Dutch Research Council (NWO) domain Applied and Engineering/ Sciences (TTW).

## References

1. Afouras, T., Owens, A., Chung, J.S., Zisserman, A.: Self-supervised learning of audio-visual objects from video. In: European Conference on Computer Vision (ECCV) (2020) 1
2. Ahsan, U., Madhok, R., Essa, I.: Video jigsaw: Unsupervised learning of spatiotemporal context for video action recognition. In: Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 179–189. IEEE (2019) 4
3. Alwassel, H., Mahajan, D., Korbar, B., Torresani, L., Ghanem, B., Tran, D.: Self-supervised learning by cross-modal audio-video clustering. In: Advances in Neural Information Processing Systems (NeurIPS). vol. 33, pp. 9758–9770 (2020) 5
4. Asano, Y.M., Patrick, M., Rupprecht, C., Vedaldi, A.: Labelling unlabelled videos from scratch with multi-modal self-supervision. In: Advances in Neural Information Processing Systems (NeurIPS) (2020) 1, 2, 5
5. Asano, Y.M., Rupprecht, C., Vedaldi, A.: A critical analysis of self-supervision, or what we can learn from a single image. In: International Conference on Learning Representations (ICLR) (2020) 2
6. Bai, Y., Fan, H., Misra, I., Venkatesh, G., Lu, Y., Zhou, Y., Yu, Q., Chandra, V., Yuille, A.: Can temporal information help with contrastive self-supervised learning? arXiv preprint arXiv:2011.13046 (2020) 5
7. 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 (CVPR). pp. 9922–9931 (2020) 4
8. Chen, B., Rouditchenko, A., Duarte, K., Kuehne, H., Thomas, S., Boggust, A., Panda, R., Kingsbury, B., Feris, R., Harwath, D., et al.: Multimodal clustering networks for self-supervised learning from unlabeled videos. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR). pp. 8012–8021 (2021) 5
9. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: Proceedings of the International Conference on Machine Learning (PMLR) (2020) 2
10. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297 (2020) 5
11. Cho, H., Kim, T., Chang, H.J., Hwang, W.: Self-supervised spatio-temporal representation learning using variable playback speed prediction. IEEE Access 9, 79562–79571 (2021) 4
12. Cole, E., Yang, X., Wilber, K., Mac Aodha, O., Belongie, S.: When does contrastive visual representation learning work? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022) 2
13. Damen, D., Doughty, H., Farinella, G.M., Furnari, A., Ma, J., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W., Wray, M.: Rescaling egocentric vision: Collection, pipeline and challenges for EPIC-KITCHENS-100. International Journal of Computer Vision (IJCV) (2021) 2, 4, 14
14. Dave, I., Gupta, R., Rizve, M.N., Shah, M.: Tclr: Temporal contrastive learning for video representation. In Computer Vision and Image Understanding (CVIU) p. 103406 (2022) 2, 5
15. Diba, A., Sharma, V., Safdari, R., Lotfi, D., Sarfraz, S., Stiefelwagen, R., Van Gool, L.: Vi2clr: Video and image for visual contrastive learning of representation.

- In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 1502–1512 (2021) 5
16. Ericsson, L., Gouk, H., Hospedales, T.M.: How well do self-supervised models transfer? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5414–5423 (2021) 2
  17. Ericsson, L., Gouk, H., Hospedales, T.M.: Why do self-supervised models transfer? investigating the impact of invariance on downstream tasks. arXiv preprint arXiv:2111.11398 (2021) 2
  18. Eyjolfssdottir, E., Branson, S., Burgos-Artizzu, X.P., Hoopfer, E.D., Schor, J., Anderson, D.J., Perona, P.: Detecting social actions of fruit flies. In: European Conference on Computer Vision. pp. 772–787 (2014) 14
  19. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 6201–6210 (2019) 11
  20. Feichtenhofer, C., Fan, H., Xiong, B., Girshick, R., He, K.: A large-scale study on unsupervised spatiotemporal representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3299–3309 (2021) 2
  21. 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 (CVPR). pp. 3636–3645 (2017) 4
  22. Gavriilyuk, K., Jain, M., Karmanov, I., Snoek, C.G.M.: Motion-augmented self-training for video recognition at smaller scale. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10429–10438 (2021) 2
  23. Ghodrati, A., Gavves, E., Snoek, C.G.M.: Video time: Properties, encoders and evaluation. In: British Machine Vision Conference (BMVC) (2018) 10
  24. Goyal, P., Mahajan, D., Gupta, A., Misra, I.: Scaling and benchmarking self-supervised visual representation learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 6391–6400 (2019) 2
  25. Goyal, R., Ebrahimi Kahou, S., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haenel, V., Freund, 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 (ICCV). pp. 5842–5850 (2017) 2, 4
  26. Grauman, K., et al.: Ego4d: Around the World in 3,000 Hours of Egocentric Video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022) 14
  27. Gu, C., Sun, C., Ross, D.A., Vondrick, C., Pantofaru, C., Li, Y., Vijayanarasimhan, S., Toderici, G., Ricco, S., Sukthankar, R., Schmid, C., Malik, J.: Ava: A video dataset of spatio-temporally localized atomic visual actions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) 2, 4, 11
  28. Han, T., Xie, W., Zisserman, A.: Video representation learning by dense predictive coding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (2019) 5
  29. Han, T., Xie, W., Zisserman, A.: Self-supervised co-training for video representation learning. In: Advances in Neural Information Processing Systems (NeurIPS) (2020) 5
  30. Hu, D., Nie, F., Li, X.: Deep multimodal clustering for unsupervised audiovisual learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9248–9257 (2019) 5



31. Huang, D., Wu, W., Hu, W., Liu, X., He, D., Wu, Z., Wu, X., Tan, M., Ding, E.: Ascnet: Self-supervised video representation learning with appearance-speed consistency. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 8096–8105 (2021) 2, 5
32. Huo, Y., Ding, M., Lu, H., Lu, Z., Xiang, T., Wen, J.R., Huang, Z., Jiang, J., Zhang, S., Tang, M., Huang, S., Luo, P.: Self-supervised video representation learning with constrained spatiotemporal jigsaw. In: Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI) (2021) 4
33. Islam, A., Chen, C.F.R., Panda, R., Karlinsky, L., Radke, R., Feris, R.: A broad study on the transferability of visual representations with contrastive learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 8845–8855 (2021) 2
34. Jenni, S., Meishvili, G., Favaro, P.: Video representation learning by recognizing temporal transformations. In: European Conference on Computer Vision (ECCV). pp. 425–442 (2020) 4
35. Jing, L., Yang, X., Liu, J., Tian, Y.: Self-supervised spatiotemporal feature learning via video rotation prediction. arXiv preprint arXiv:1811.11387 (2018) 4
36. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., Suleyman, M., Zisserman, A.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017) 1, 2, 5
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. pp. 8545–8552 (2019) 4
38. Kolesnikov, A., Zhai, X., Beyer, L.: Revisiting self-supervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1920–1929 (2019) 2
39. Köpüklü, O., Wei, X., Rigoll, G.: You only watch once: A unified CNN architecture for real-time spatiotemporal action localization. arXiv preprint arXiv:1911.06644 (2019) 10
40. Korbar, B., Tran, D., Torresani, L.: Cooperative learning of audio and video models from self-supervised synchronization. In: Advances in Neural Information Processing Systems (NeurIPS). vol. 31 (2018) 5
41. Kotar, K., Ilharco, G., Schmidt, L., Ehsani, K., Mottaghi, R.: Contrasting contrastive self-supervised representation learning pipelines. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9949–9959 (2021) 2
42. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: HMDB: a large video database for human motion recognition. In: Proceedings of the International Conference on Computer Vision (ICCV) (2011) 1, 2
43. Li, Y., Li, Y., Vasconcelos, N.: Resound: Towards action recognition without representation bias. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 513–528 (2018) 2
44. Lin, Y., Guo, X., Lu, Y.: Self-supervised video representation learning with meta-contrastive network. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 8239–8249 (2021) 2, 5
45. Luo, D., Liu, C., Zhou, Y., Yang, D., Ma, C., Ye, Q., Wang, W.: Video cloze procedure for self-supervised spatio-temporal learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 11701–11708 (2020) 4

46. Ma, S., Zeng, Z., McDuff, D., Song, Y.: Active contrastive learning of audio-visual video representations. In: International Conference on Learning Representations (ICLR) (2021) 5
47. Mettes, P., Gemert, J.C.v., Snoek, C.G.M.: Spot on: Action localization from pointly-supervised proposals. In: European conference on computer vision (ECCV). pp. 437–453. Springer (2016) 10
48. Misra, I., Zitnick, C.L., Hebert, M.: Shuffle and learn: unsupervised learning using temporal order verification. In: European Conference on Computer Vision (ECCV). pp. 527–544 (2016) 1, 4
49. Morgado, P., Vasconcelos, N., Misra, I.: Audio-visual instance discrimination with cross-modal agreement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 1, 2, 5
50. Newell, A., Deng, J.: How useful is self-supervised pretraining for visual tasks? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) 2
51. Ng, X.L., Ong, K.E., Zheng, Q., Ni, Y., Yeo, S.Y., Liu, J.: Animal kingdom: A large and diverse dataset for animal behavior understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 19023–19034 (2022) 14
52. Nguyen, T., Raghu, M., Kornblith, S.: Do wide and deep networks learn the same things? uncovering how neural network representations vary with width and depth. In: International Conference on Learning Representations (ICLR) (2021) 13, 14
53. Pan, T., Song, Y., Yang, T., Jiang, W., Liu, W.: Videomoco: Contrastive video representation learning with temporally adversarial examples. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11205–11214 (2021) 1, 5
54. Patrick, M., Asano, Y.M., Kuznetsova, P., Fong, R., Henriques, J.F., Zweig, G., Vedaldi, A.: Multi-modal self-supervision from generalized data transformations. In: International Conference on Computer Vision (ICCV) (2021) 1, 2, 5
55. Pedersen, M., Haurum, J.B., Bengtson, S.H., Moeslund, T.B.: 3d-zef: A 3d zebrafish tracking benchmark dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2426–2436 (2020) 14
56. Peihao, C., Deng, H., Dongliang, H., Xiang, L., Runhao, Z., Shilei, W., Mingkui, T., Chuang, G.: Rspnet: Relative speed perception for unsupervised video representation learning. In: The AAAI Conference on Artificial Intelligence (AAAI) (2021) 1, 2, 5
57. Piergiovanni, A., Angelova, A., Ryoo, M.S.: Evolving losses for unsupervised video representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 133–142 (2020) 1
58. 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 (CVPR). pp. 6964–6974 (2021) 1, 2, 5, 6
59. Recasens, A., Luc, P., Alayrac, J.B., Wang, L., Strub, F., Tallec, C., Malinowski, M., Pătrăucean, V., Altché, F., Valko, M., et al.: Broaden your views for self-supervised video learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 1255–1265 (2021) 2, 6
60. Sariyildiz, M.B., Kalantidis, Y., Larlus, D., Alahari, K.: Concept generalization in visual representation learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9629–9639 (2021) 2

61. Schiappa, M.C., Rawat, Y.S., Shah, M.: Self-supervised learning for videos: A survey. arXiv preprint arXiv:2207.00419 (2022) 5
62. Shahroudy, A., Liu, J., Ng, T.T., Wang, G.: Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1010–1019 (2016) 4
63. Shao, D., Zhao, Y., Dai, B., Lin, D.: Finegym: A hierarchical video dataset for fine-grained action understanding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020) 4, 8, 9
64. Sigurdsson, G.A., Varol, G., Wang, X., Farhadi, A., Laptev, I., Gupta, A.: Hollywood in Homes: Crowdsourcing Data Collection for Activity Understanding. In: European Conference on Computer Vision (ECCV). pp. 510 – 526 (2016) 4, 11
65. 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) 1, 2, 4
66. Sun, C., Nagrani, A., Tian, Y., Schmid, C.: Composable augmentation encoding for video representation learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 8834–8844 (2021) 5
67. Sun, J.J., Karigo, T., Chakraborty, D., Mohanty, S., Wild, B., Sun, Q., Chen, C., Anderson, D., Perona, P., Yue, Y., Kennedy, A.: The multi-agent behavior dataset: Mouse dyadic social interactions. In: Vanschoren, J., Yeung, S. (eds.) Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (2021) 14
68. Suzuki, T., Itazuri, T., Hara, K., Kataoka, H.: Learning spatiotemporal 3d convolution with video order self-supervision. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops. pp. 0–0 (2018) 4
69. 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 (ACM MM). pp. 2193–2201 (2020) 5
70. Tao, L., Wang, X., Yamasaki, T.: Pretext-contrastive learning: Toward good practices in self-supervised video representation learning. arXiv preprint arXiv:2010.15464 (2021) 2, 5
71. Thoker, F.M., Doughty, H., Snoek, C.: Skeleton-contrastive 3d action representation learning. In: in Proceedings of the 29th ACM International Conference on Multimedia,(ACM MM ) (2021) 5
72. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6450–6459 (2018) 5
73. Van Horn, G., Cole, E., Beery, S., Wilber, K., Belongie, S., Mac Aodha, O.: Benchmarking representation learning for natural world image collections. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12884–12893 (2021) 2
74. Wallace, B., Hariharan, B.: Extending and analyzing self-supervised learning across domains. In: European Conference on Computer Vision (ECCV). pp. 717–734. Springer (2020) 2
75. Wang, G., Zhou, Y., Luo, C., Xie, W., Zeng, W., Xiong, Z.: Unsupervised visual representation learning by tracking patches in video. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 1, 2, 4, 5
76. Wang, J., Jiao, J., Bao, L., He, S., Liu, Y., Liu, W.: Self-supervised spatio-temporal representation learning for videos by predicting motion and appearance statistics.

- In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4006–4015 (2019) 4
77. Wang, J., Jiao, J., Liu, Y.H.: Self-supervised video representation learning by pace prediction. In: European Conference on Computer Vision (ECCV). pp. 504–521 (2020) 4
  78. Wang, J., Gao, Y., Li, K., Lin, Y., Ma, A.J., Cheng, H., Peng, P., Ji, R., Sun, X.: Removing the background by adding the background: Towards background robust self-supervised video representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 2
  79. Wei, D., Lim, J.J., Zisserman, A., Freeman, W.T.: Learning and using the arrow of time. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8052–8060 (2018) 10
  80. Xiao, F., Tighe, J., Modolo, D.: Modist: Motion distillation for self-supervised video representation learning. arXiv preprint arXiv:2106.09703 (2021) 2
  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 (CVPR). pp. 10334–10343 (2019) 1, 4
  82. Yang, C., Xu, Y., Dai, B., Zhou, B.: Video representation learning with visual tempo consistency. arXiv preprint arXiv:2006.15489 (2020) 2, 5, 6
  83. Yang, X., He, X., Liang, Y., Yang, Y., Zhang, S., Xie, P.: Transfer learning or self-supervised learning? a tale of two pretraining paradigms. arXiv preprint arXiv:2007.04234 (2020) 2
  84. Yao, T., Zhang, Y., Qiu, Z., Pan, Y., Mei, T.: Seco: Exploring sequence supervision for unsupervised representation learning. In: AAAI. vol. 2, p. 7 (2021) 4
  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 (CVPR). pp. 6548–6557 (2020) 4
  86. Zhai, X., Puigcerver, J., Kolesnikov, A., Ruysen, P., Riquelme, C., Lucic, M., Djolonga, J., Pinto, A.S., Neumann, M., Dosovitskiy, A., et al.: A large-scale study of representation learning with the visual task adaptation benchmark. arXiv preprint arXiv:1910.04867 (2019) 2
  87. 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) 10
  88. Zhang, Y., Po, L.M., Xu, X., Liu, M., Wang, Y., Ou, W., Zhao, Y., Yu, W.Y.: Contrastive spatio-temporal pretext learning for self-supervised video representation. In: Proceedings of the AAAI Conference on Artificial Intelligence (2022) 5