Semi-Supervised Action Recognition with Temporal Contrastive Learning

Ankit Singh^{1*} Omprakash Chakraborty^{2*} Ashutosh Varshney² Rameswar Panda³ Rogerio Feris³ Kate Saenko^{3,4} Abir Das² ¹ IIT Madras, ² IIT Kharagpur, ³ MIT-IBM Watson AI Lab, ⁴ Boston University

Abstract

Learning to recognize actions from only a handful of labeled videos is a challenging problem due to the scarcity of tediously collected activity labels. We approach this problem by learning a two-pathway temporal contrastive model using unlabeled videos at two different speeds leveraging the fact that changing video speed does not change an action. Specifically, we propose to maximize the similarity between encoded representations of the same video at two different speeds as well as minimize the similarity between different videos played at different speeds. This way we use the rich supervisory information in terms of 'time' that is present in otherwise unsupervised pool of videos. With this simple yet effective strategy of manipulating video playback rates, we considerably outperform video extensions of sophisticated state-of-the-art semi-supervised image recognition methods across multiple diverse benchmark datasets and network architectures. Interestingly, our proposed approach benefits from out-of-domain unlabeled videos showing generalization and robustness. We also perform rigorous ablations and analysis to validate our approach. Project page: https://cvir.github.io/TCL/.

1. Introduction

Supervised deep learning approaches have shown remarkable progress in video action recognition [7, 15, 16, 17, 36, 49]. However, being supervised, these models are critically dependent on large datasets requiring tedious human annotation effort. This motivates us to look beyond the supervised setting as supervised methods alone may not be enough to deal with the volume of information contained in videos. Semi-supervised learning approaches use structural invariance between different views of the same data as a source of supervision for learning useful representations. In recent times, semi-supervised representation learning models [10, 29, 38, 50] have performed very well even surpassing its supervised counterparts in case of images [22, 47].



Figure 1: Comparison of top-1 accuracy for TCL (Ours) with Pseudo-Label [35] and FixMatch [47] baselines trained with different percentages of labeled training data. We evaluate the efficacy of the approaches in terms of the least proportion of labeled data required to surpass the fully supervised [36] performance (shown with the red dotted line). With only 33% and 15% of labeled data, our proposed TCL framework surpasses the supervised approaches in Mini-Something-V2 [23] and Jester [37] datasets respectively. The two other compared methods fail to reach the accuracy of the fully supervised approach with such small amount of labeled data. (Best viewed in color.)

Notwithstanding their potential, semi-supervised video action recognition has received very little attention. Trivially extending the image domain approaches to videos without considering the rich temporal information may not quite bridge the performance gap between the semi and the fully supervised learning. But, in videos, we have another source of supervision: time. We all know that an action recognizer is good if it can recognize actions irrespective of whether the actions are performed slowly or quickly. Recently, supervised action recognition has benefited a lot by using differently paced versions of the same video during training [17, 54]. Motivated by the success of using slow and fast versions of videos for supervised action recognition as well as by the success of the contrastive learning frameworks [26, 41], we propose Temporal Contrastive Learning (TCL) for semi-supervised action recognition in videos where consistent features representing both slow and fast versions of the same videos are learned.

Starting with a model trained with limited labeled data, we present a two-pathway model that processes unlabeled videos at two different speeds and finds their representations. Though played at two different speeds, the videos share the same semantics. Thus, similarity between these

^{*}The first two authors contributed equally.

representations are maximized. Likewise, the similarity between the representations of different videos are minimized. We achieve this by minimizing a modified NT-Xent contrastive loss [10, 50] between these videos with different playback rates. While minimizing a contrastive loss helps to produce better visual representations by learning to be invariant to different views of the data, it ignores information shared among samples of same action class as the loss treats each video individually. To this end, we propose a new perspective of contrastive loss between neighborhoods. Neighborhoods are compact groups of unlabeled videos with high class consistency. In absence of ground-truth labels, groups are formed by clustering videos with same pseudo-labels and are represented by averaging the representations of the constituent videos. Contrastive objective between groups formed off the two paths explores the underlying class concept that traditional NT-Xent loss among individual video instances does not take into account. We term the contrastive loss considering only individual instances as the instance-contrastive loss and the same between the groups as the group-contrastive loss respectively.

We perform extensive experiments on four standard datasets and demonstrate that TCL achieves superior performance over extended baselines of state-of-the-art image domain semi-supervised approaches. Figure 1 shows comparison of **TCL** with Pseudo-Label [35] and FixMatch [47] trained using different percentages of labeled training data. Using the same backbone network (ResNet-18), TCL needs only 33% and 15% of labeled data in Mini-Something-V2 [9] and Jester [37] respectively to reach the performance of the fully supervised approach [36] that uses 100% labeled data. On the other hand, the two compared methods fail to reach the accuracy of the fully supervised approach with such small amount of labeled data. Likewise, we observe as good as 8.14% and 4.63% absolute improvement in recognition performance over the next best approach, Fix-Match [47] using only 5% labeled data in Mini-Something-V2 [9] and Kinetics-400 [32] datasets respectively. In a new realistic setting, we argue that unlabeled videos may come from a related but different domain than that of the labeled data. For instance, given a small set of labeled videos from a third person view, our approach is shown to benefit from using only first person unlabeled videos on Charades-Ego [44] dataset, demonstrating the robustness to domain shift in the unlabeled set. To summarize, our key contributions include:

- First of all, we treat the time axis in unlabeled videos specially, by processing them at two different speeds and propose a two-pathway temporal contrastive semi-supervised action recognition framework.
- Next, we identify that directly employing contrastive objective instance-wise on video representations learned with different frame-rates may miss crucial information shared across samples of same in-

herent class. A novel group-contrastive loss is pioneered to couple discriminative motion representation with pace-invariance that significantly improves semisupervised action recognition performance.

 We demonstrate through experimental results on four datasets, TCL's superiority over extended baselines of successful image-domain semi-supervised approaches. The versatility and robustness of our approach in case of training with unlabeled videos from a different domain is shown along with in-depth ablation analysis pinpointing the role of the different components.

2. Related Work

Action Recognition. Action recognition is a challenging problem with great application potential. Conventional approaches based on deep neural networks are mostly built over a two-stream CNN based framework [46], one to process a single RGB frame and the other for optical flow input to analyze the spatial and temporal information respectively. Many variants of 3D-CNNs such as C3D [49], I3D [7] and ResNet3D [27], that use 3D convolutions to model space and time jointly, have also been introduced for action recognition. SlowFast network [17] employs two pathways for recognizing actions by processing a video at both slow and fast frame rates. Recent works also utilize 2D-CNNs for efficient video classification by using different temporal aggregation modules such as temporal averaging in TSN [52], bag of features in TRN [61], channel shifting in TSM [36], depthwise convolutions in TAM [15]. Despite promising results on common benchmarks, these models are critically dependent on large datasets that require careful and tedious human annotation effort. In contrast, we propose a simple yet effective temporal contrastive learning framework for semi-supervised action recognition that alleviates the data annotation limitation of supervised methods.

Semi-Supervised Learning. Semi-supervised learning (SSL) has been studied from multiple aspects (see reviews [8]). Various strategies have been explored *e.g.*, generative models [40, 42], self-training using pseudolabels [1, 24, 35] and consistency regularization [2, 4, 5, 34, 39, 48, 55]. Leveraging self-supervised learning like rotation prediction [20] and image transformations [13] is also another recent trend for SSL [60]. While there has been tremendous progress in semi-supervised image classification, SSL for action recognition is still a novel and rarely addressed problem. Iosifidis et al. [31], first utilize traditional Action Bank for action representation and then uses a variant of extreme learning machine for semi-supervised classification. The work most related to ours is [47] which first generates confident one-hot labels for unlabelled images and then trains the model to be consistent across different forms of image augmentations. While this has re-



Figure 2: **Illustration of our Temporal Contrastive Learning (TCL) Framework.** Our approach consists of *base* and *auxiliary* pathways that share the same weights. Base pathway accepts video frames sampled at a higher rate while the auxiliary pathway takes in frames at a lower framerate. At first, the base network is trained using limited labeled data. Subsequently, the auxiliary pathway comes into picture for the unlabeled samples by encouraging video representations to match in both pathways in absence of labels. This is done by maximizing agreement between the outputs of the two pathways for a video while minimizing the same for different videos. In addition, originally unlabeled videos with high semantic similarity are grouped by pseudo-labels assigned to them. To exploit the high consistency and compactness of group members, the average representations of groups with same pseudo-label in different pathways are made similar while that between the varying groups are made maximally different. Two separate contrastive losses (ref Sections 3.2.2 and 3.2.3) are used for this purpose. Given a video at test time, only the base network is used to recognize the action. (Best viewed in color.)

cently achieved great success, the data augmentations for generating different transformations are limited to transformations in the image space and fail to leverage the temporal information present in videos. We differ from [47] as we propose a temporal contrastive learning framework for semi-supervised action recognition by modeling temporal aspects using two pathways at different speeds instead of augmenting images. We further propose a group-wise contrastive loss in addition to instance-wise contrastive loss for learning discriminative features for action recognition.

Contrastive Learning. Contrastive learning [10, 11, 19, 28, 30, 33, 38, 50, 53] is becoming increasingly attractive due to its great potential to leverage large amount of unlabeled data. The essence of contrastive learning lie in maximizing the similarity of representations among positive samples while encouraging discrimination for negative samples. Some recent works have also utilized contrastive learning [21, 25, 41, 43, 58] for self-supervised video representation learning. Spatio-temporal contrastive learning using different augmentations for learning video features is presented in [41]. Speed of a video is also investigated for self-supervised learning [3, 51, 59] unlike the problem we consider in this paper. While our approach is inspired by these, we focus on semi-supervised action recognition in videos, where our goal is to learn consistent features representing two different frame rates of the unlabeled videos.

3. Methodology

In this section, we present our novel semi-supervised approach to efficiently learn video representations. First we briefly discuss the problem description and then describe our framework and its components in detail.

3.1. Problem Setup

Our aim is to address semi-supervised action recognition where only a small set of videos (\mathcal{D}_l) has labels, but a large number of unlabeled videos (\mathcal{D}_u) are assumed to be present alongside. The set $\mathcal{D}_l \triangleq \{V^i, y^i\}_{i=1}^{N_l}$ comprises of N_l videos where the i^{th} video and the corresponding activity label is denoted by V^i and y^i respectively. For a dataset of videos with C different activities, y^i is often assumed to be an element of the label set $Y = \{1, 2, \dots, C\}$. Similarly, the unlabeled set $\mathcal{D}_u \triangleq \{U^i\}_{i=1}^{N_u}$ comprises of $N_u(\gg N_l)$ videos without any associated labels. We use the unlabeled videos at two different frame rates and refer to them as fast and slow videos. The fast version of the video U^i is represented as a collection of M frames *i.e.*, $U_f^i = \{F_{f,1}^i, F_{f,2}^i, \dots, F_{f,M}^i\}$. Likewise, the slow version of the same is represented as $U_s^i = \{F_{s,1}^i, F_{s,2}^i, \dots, F_{s,N}^i\}$ where N < M. The frames are sampled from the video following Wang *et. al* [52] where a random frame is sampled uniformly from consecutive non-overlapping segments.

3.2. Temporal Contrastive Learning

As shown in Figure 2, our 'Temporal Contrastive Learning (**TCL**)' framework processes the input videos in two pathways, namely *base* and *auxiliary* pathways. The fast version of the videos are processed by *base* while the slow versions are processed by the *auxiliary* pathway. Both pathways share same neural backbone (denoted by g(.)). Different stages of training in **TCL** framework are described next.

3.2.1 Supervised Training Stage

The neural network backbone is initially trained using only the small labeled data \mathcal{D}_l by passing it through the base branch. Depending on whether the backbone involves 2D [36, 52] or 3D convolution [7, 27] operations, the representation $(g(V^i))$ of the video V^i used in our framework is average of the frame logits or the logits from the 3D backbone respectively. We minimize the standard supervised cross-entropy loss (\mathcal{L}_{sup}) on the labeled data as follows.

$$\mathcal{L}_{sup} = -\sum_{c=1}^{C} (y^i)_c \log(g(V^i))_c \tag{1}$$

3.2.2 Instance-Contrastive Loss

1

Equipped with an initial backbone trained with limited supervision, our goal is to learn a model that can use a large pool of unlabeled videos for better activity understanding. To this end, we use temporal co-occurrence of unlabeled activities at multiple speeds as a proxy task and enforce this with a pairwise contrastive loss. Specifically, we adjust the frame sampling rate to get videos with different speeds. Let us consider a minibatch with B unlabeled videos. The model is then trained to match the representation $g(U_f^i)$ of the comparatively faster version of the video (U^i) with that $(g(U^i_s))$ of the slower version. $g(U^i_f)$ and $g(U^i_s)$ forms the positive pair. For rest of B-1 videos, $g(U_f^i)$ and $g(U_p^k)$ form negative pairs where representation of k^{th} video can come from either of the pathways (*i.e.*, $p \in \{f, s\}$). As different videos forming the negative pairs, have different content, the representation of different videos in either of the pathways are pushed apart. This is achieved by employing a contrastive loss (\mathcal{L}_{ic}) as follows.

$$\mathcal{L}_{ic}(U_f^i, U_s^i) = -\log \frac{h(g(U_f^i), g(U_s^i))}{h(g(U_f^i), g(U_s^i)) + \sum_{\substack{k=1\\p \in \{s, f\}}}^B \mathbb{1}_{\{k \neq i\}} h(g(U_f^i), g(U_p^k))}$$

$$(2)$$

where $h(\mathbf{u}, \mathbf{v}) = \exp\left(\frac{\mathbf{u}^{\top}\mathbf{v}}{||\mathbf{u}||_2||\mathbf{v}||_2}/\tau\right)$ is the exponential of cosine similarity measure and τ is the temperature hyperparameter. The final instance-contrastive loss is computed for all positive pairs, i.e., both (U_f^i, U_s^i) and (U_s^i, U_f^i) across minibatch. The loss function encourages to decrease the similarity not only between different videos in two pathways but also between different videos across both of them.



Figure 3: Advantage of group-contrastive loss over instancecontrastive loss. A contrastive objective between instances may try to push different instances of same action apart (right), while forming groups of videos with same activity class avoids such inadvertent competition (left). In absence of true labels, such grouping is done by the predicted pseudo-labels. (Best viewed in color.)

3.2.3 Group-Contrastive Loss

Directly applying contrastive loss between different video instances in absence of class-labels does not take the high level action semantics into account. As illustrated in Figure 3, such a strategy can inadvertently learn different representations for videos containing same actions. We employ contrastive loss among groups of videos with similar actions where relations within the neighborhood of different videos are explored. Specifically, each unlabeled video U^i in each of the two pathways are assigned pseudo-labels that correspond to the class having the maximum activation. Let \hat{y}_f^i and \hat{y}_s^i denote the pseudo-labels of the videos U^i in the fast and the slow pathways respectively. Videos having the same pseudo-label in a minibatch form a group in each pathway and the average of the representations of constituent videos

$$R_{p}^{l} = \frac{\sum_{i=1}^{B} \mathbb{1}_{\{\hat{y}_{p}^{i}=l\}}g(U_{p}^{i})}{T}$$
(3)

where $\mathbb{1}$ is an indicator function that evaluates to 1 for the videos with pseudo-label equal to $l \in Y$ in each pathway $p \in \{f, s\}$. T is the number of such videos in the minibatch.

Considering the high class consistency among two groups with same label in two pathways, we require these groups to give similar representations in the feature space. Thus, in the group-contrastive objective, all pairs (R_f^l, R_s^l) act as positive pairs while the negative pairs are the pairs (R_f^l, R_p^m) with $p \in \{f, s\}$ and $m \in Y \setminus l$ such that the constituent groups are different in either of the pathways. The group-contrastive loss involving these pairs is,

$$\mathcal{L}_{gc}(R_{f}^{l}, R_{s}^{l}) = -\log \frac{h(R_{f}^{l}, R_{s}^{l})}{h(R_{f}^{l}, R_{s}^{l}) + \sum_{\substack{m=1\\p \in \{s, f\}}}^{C} \mathbb{1}_{\{m \neq l\}} h(R_{f}^{l}, R_{p}^{m})}$$
(4)

Similar to instance-contrastive loss, group-contrastive loss is also computed for all positive pairs - both (R_f^l, R_s^l)

and (R_s^l, R_f^l) across the minibatch. Overall, the loss function for training our model involving the limited labeled data and the unlabeled data is,

$$\mathcal{L} = \mathcal{L}_{sup} + \gamma * \mathcal{L}_{ic} + \beta * \mathcal{L}_{gc}$$
⁽⁵⁾

where, γ and β are weights of the instance-contrastive and group-contrastive losses respectively.

3.3. TCL with Pretraining and Finetuning

Self-supervised pretraining has recently emerged as a promising alternative, which not only avoids huge annotation effort but also is better and robust compared to its supervised counterpart in many visual tasks [14, 56, 62]. Motivated by this, we adopt self-supervised pretraining to initialize our model with very minimal change in the framework. Specifically, we employ self-supervised pretraining at the beginning by considering the whole of the labeled and the unlabeled data $\mathcal{D}_l \cup \mathcal{D}_u$ as unlabeled data only and using instance-contrastive loss \mathcal{L}_{ic} to encourage consistency between representations learned in the two pathways (ref. Eq. 2). These weights are then used to initialize the base and the auxiliary pathways before our approach commences for semi-supervised learning of video representations. For effective utilization of unlabeled data, we also finetune the base pathway with pseudo-labels [35] generated at the end of our contrastive learning, which greatly enhances the discriminability of the features, leading to improvement in recognition performance. We empirically show that starting with the same amount of labeling, both selfsupervised pretraining and finetuning with pseudo-labels (Pretraining \rightarrow **TCL** \rightarrow Finetuning) benefits more compared to the same after limited supervised training only.

4. Experiments

In this section, we conduct extensive experiments to show that our **TCL** framework outperforms many strong baselines on several benchmarks including one with domain shift. We also perform comprehensive ablation experiments to verify the effectiveness of different components in detail.

4.1. Experimental Setup

Datasets. We evaluate our approach using four datasets, namely Mini-Something-V2 [9], Jester [37], Kinetics-400 [32] and Charades-Ego [44]. Mini-Something-V2 is a subset of Something-Something V2 dataset [23] containing 81K training videos and 12K testing videos across 87 action classes. Jester [37] contains 119K videos for training and 15K videos for validation across 27 annotated classes for hand gestures. Kinetics-400 [32] is one of the most popular large-scale benchmark for video action recognition. It consists of 240K videos for training and 20K videos for validation across 400 action categories, with each video lasting 6-10 seconds. Charades-Ego [44] contains 7,860 untrimmed

egocentric videos of daily indoors activities recorded from both third and first person views. The dataset contains 68,536 temporal annotations for 157 action classes. We use a subset of the third person videos from Charades-Ego as the labeled data while the first person videos are considered as unlabeled data to show the effectiveness of our approach under domain shift in the unlabeled data. More details about the datasets are included in supplementary material.

Baselines. We compare our approach with the following baselines and existing semi-supervised approaches from 2D image domain extended to video data. First, we consider a supervised baseline where we train an action classifier having the same architecture as the base pathway of our approach. This is trained using a small portion of the labeled examples assuming only a small subset of labeled examples is available as annotated data. Second, we compare with state-of-the-art semi-supervised learning approaches, including Pseudo-Label [35] (ICMLW'13), Mean Teacher [48] (NeurIPS'17), S4L [60] (ICCV'19), Mix-Match [5] (NeurIPS'19), and FixMatch [47] (NeurIPS'20). We use same backbone and experimental settings for all the baselines (including our approach) for a fair comparison.

Implementation Details. We use Temporal Shift Module (TSM) [36] with ResNet-18 backbone as the base action classifier in all our experiments. We further investigate performance of different methods by using ResNet-50 on Mini-Something-V2 dataset. TSM has recently shown to be very effective due to its hardware efficiency and lesser computational complexity. We use uniformly sampled 8 and 4 frame segments from unlabeled videos as input to the base and the auxiliary pathways respectively to process unlabeled videos in our TCL framework. On the other hand, we use only 8 frame segments for labeled videos and compute the final performance using 8 frame segments in the base pathway for all the methods. Note that our approach is agnostic to the backbone architecture and particular values of frame rates. Following the standard practice [47] in SSL, we randomly choose a certain percentage of labeled samples as a small labeled set and discard the labels for the remaining data to form a large unlabeled set. Our approach is trained with different percentages of labeled samples for each dataset (1%, 5% and 10%). We train our models for 400 epochs where we first train our model with supervised loss \mathcal{L}_{sup} using only labeled data for 50 epochs. We then train our model using the combined loss (ref. Eq. 5) for the next 300 epochs. Finally, for finetuning with pseudo-labels, we train our model with both labeled and unlabeled videos having pseudo-label confidence more than 0.8 for 50 epochs.

During pretraining, we follow the standard practice in self-supervised learning [10, 58] and train our model using all the training videos without any labels for 200 epochs. We use SGD [6] with a learning rate of 0.02 and momentum value of 0.9 with cosine learning rate decay in all our experi-

	ResNet-18			ResNet-50		
Approach	1%	5%	10%	1%	5%	10%
Supervised (8f)	5.98 ± 0.68	17.26 ± 1.17	24.67 ± 0.68	5.69 ± 0.51	$16.68 {\pm} 0.25$	$25.92{\pm}0.53$
Pseudo-Label [35] (ICMLW'13)	6.46 ± 0.32	$18.76 {\pm} 0.77$	$25.67 {\pm} 0.45$	$6.66 {\pm} 0.89$	$18.77 {\pm} 1.18$	$28.85{\pm}0.91$
Mean Teacher [48] (NeurIPS' 17)	7.33 ± 1.13	$20.23 {\pm} 1.59$	$30.15 {\pm} 0.42$	6.82 ± 0.18	$21.80{\pm}1.54$	$32.12 {\pm} 2.37$
S4L [60] (ICCV' 19)	7.18 ± 0.97	$18.58 {\pm} 1.05$	$26.04{\pm}1.89$	6.87 ± 1.29	$17.73 {\pm} 0.26$	$27.84 {\pm} 0.75$
MixMatch [5] (NeurIPS'19)	7.45 ± 1.01	$18.63 {\pm} 0.99$	$25.78 {\pm} 1.01$	6.48 ± 0.83	$17.77 {\pm} 0.12$	$27.03 {\pm} 1.66$
FixMatch [47] (NeurIPS' 20)	6.04 ± 0.44	$21.67{\pm}0.18$	$33.38{\pm}1.58$	$6.54{\pm}0.71$	$25.34{\pm}2.03$	$37.44{\pm}1.31$
TCL (Ours)	7.79 ± 0.57	$29.81 {\pm} 0.77$	38.61 ± 0.91	7.54 ± 0.32	27.22 ± 1.86	$40.70 {\pm} 0.42$
TCL w/ Finetuning	8.65 ± 0.76	$30.55 {\pm} 1.36$	40.06 ± 1.14	$8.56 {\pm} 0.31$	$28.84{\pm}1.22$	$41.68 {\pm} 0.56$
TCL w/ Pretraining & Finetuning	9.91±1.84	$30.97{\pm}0.07$	$41.55{\pm}0.47$	9.19 ± 0.43	$29.85{\pm}1.76$	$41.33{\pm}1.07$

Table 1: Performance Comparison in Mini-Something-V2. Numbers show average Top-1 accuracy values with standard deviations over 3 random trials for different percentages of labeled data. **TCL** significantly outperforms all the compared methods in both cases.

	Jester			Kinetics-400	
Approach	1%	5%	10%	1%	5%
Supervised (8f)	52.55 ± 4.36	85.22 ± 0.61	90.45 ± 0.33	6.17 ± 0.32	20.50 ± 0.23
Pseudo-Label [35] (ICMLW'13)	57.99±3.70	$87.47 {\pm} 0.64$	90.96±0.48	6.32 ± 0.19	$20.81 {\pm} 0.86$
Mean Teacher [48] (NeurIPS'17)	56.68 ± 1.46	$88.80 {\pm} 0.44$	92.07±0.03	$6.80 {\pm} 0.42$	$22.98 {\pm} 0.43$
S4L [60] (ICCV'19)	64.98 ± 2.70	$87.23 {\pm} 0.15$	90.81±0.32	$6.32 {\pm} 0.38$	$23.33 {\pm} 0.89$
MixMatch [5] (NeurIPS'19)	58.46 ± 3.26	$89.09 {\pm} 0.21$	92.06±0.46	$6.97 {\pm} 0.48$	$21.89 {\pm} 0.22$
FixMatch [47] (NeurIPS'20)	61.50 ± 0.77	$90.20{\pm}0.35$	92.62±0.60	$6.38{\pm}0.38$	$25.65{\pm}0.28$
TCL (Ours)	75.21±4.48	$93.29 {\pm} 0.24$	94.64±0.21	7.69 ± 0.21	30.28 ± 0.13
TCL w/ Finetuning	77.25 ± 4.02	$93.53 {\pm} 0.15$	94.74 ± 0.25	$8.45 {\pm} 0.25$	$31.50 {\pm} 0.23$
TCL w/ Pretraining & Finetuning	82.55±1.94	$93.73{\pm}0.25$	94.93±0.02	11.56 ± 0.22	$31.91{\pm}0.46$

Table 2: **Performance Comparison on Jester and Kinetics-400**. Numbers show the top-1 accuracy values using ResNet-18 on both datasets. Our approach **TCL** achieves the best performance across different percentages of labeled data.

ments. Given a mini-batch of labeled samples B_l , we utilize $\mu \times B_l$ unlabeled samples for training. We set μ to 3 and τ to 0.5 in all our experiments. γ and β values are taken to be 9 and 1 respectively, unless otherwise mentioned. Random scaling and cropping are used as data augmentation during training (and we further adopt random flipping for Kinetics-400), as in [36]. Following [36], we use just 1 clip per video and the center 224×224 crop for evaluation. More implementation details are included in supplementary material.

4.2. Large-scale Experiments and Comparisons

Tables 1- 3 show performance of different methods on all four datasets, in terms of average top-1 clip accuracy and standard deviation over 3 random trials.

Mini-Something-V2. Table 1 shows the performance comparison with both ResNet-18 (left half) and ResNet-50 (right half) backbones on Mini-Something-V2. **TCL** outperforms the video extensions of all the semi-supervised image-domain baselines for all three percentages of labeled training data. The improvement is especially prominent for low capacity model (ResNet-18) and low data (only 1% and 5% data with labels) regime. Notably, our approach outperforms the most recent approach, FixMatch by 1.75% while training with only 1% labeled data. The improvement is 8.14% for the case when 5% data is labeled. These improvements clearly show that our approach is able to leverage the temporal information more effectively compared to FixMatch that focuses on only spatial image augmentations. Figure 4 shows the classwise improvement over FixMatch along with the number of labeled training data per class in the case of 5% labeling. The plot shows that a overwhelming majority of the activities experienced improvement with decrease in performance for only 1 class out of 18 having less than 20 labeled videos per class (right of the figure). For low labeled-data regime (1% and 5%), a heavier ResNet-50 model shows signs of overfitting as is shown by slight drop in performance. On the other hand, using ResNet-50 backbone instead of ResNet-18 is shown to benefit TCL if the model is fed with more labeled data. Moreover, TCL with finetuning and pretraining shows further improvement, leading to best performance in both cases.

Jester. Our approach TCL also surpasses the performance of existing semi-supervised approaches in Jester as shown in Table 2 (left). In particular, TCL achieves 10.23% absolute improvement compared to S4L (the next best) in very low labeled-data regime (1% only). Adding finetuning and self-supervised pretraining further increases this difference



Figure 4: Change in classwise top-1 accuracy of TCL over Fix-Match on Mini-Something-V2. Blue bars show the change in accuracy on 5% labeled scenario, while the red line shows the number of labeled videos per class (sorted). Compared to Fix-Match, TCL improves the performance of most classes including those with less labeled data. (Best viewed in color.)

to 17.57%. Furthermore, **TCL** with pretraining and finetuning achieves a top-1 accuracy of 94.93% using 10% labeled data which is only 0.32% lower than the fully supervised baseline trained using all the labels (95.25%).

Kinetics-400. Table 2 (right) summarizes the results on Kinetics-400, which is one of the widely used action recognition datasets consisting of 240K videos across 400 classes. **TCL** outperforms FixMatch by a margin of 1.31% and 4.63% on 1% and 5% scenarios respectively, showing the superiority of our approach on large scale datasets. The top-1 accuracy achieved using **TCL** with finetuning and pre-training is almost twice better than the supervised approach when only 1% of the labeled data is used. The results also show that off-the-shelf extensions of sophisticated state-of-the-art semi-supervised image classification methods offer little benefit to action classification on videos.

Charades-Ego. We use third person videos from Charades-Ego [44] as the target while first person videos form the additional unlabeled set. During training, labeled data is taken only from the target domain while unlabeled data is obtained from both the target and the domain-shifted videos. To modulate domain shift in unlabeled data, we introduce a new hyperparameter ρ , whose value denotes the proportion of target videos in the unlabeled set. For a fixed number of unlabeled videos $|\mathcal{D}_u|$, we randomly select $\rho \times |D_u|$ videos from the target while the remaining $(1 - \rho) \times |D_u|$ are selected from the other domain. Following the standard practice [12] in this dataset, we first pretrain the model using Charades [45] and experimented using three different values of ρ : 1,0.5,0 for 10% target data with labels. Table 3 shows the mean Average Precision (mAP) of our method including the supervised approach, PseudoLabel and Fix-Match. TCL outperforms both methods by around 1% mAP for all three ρ values. In the case when all the unlabeled data is from the shifted domain ($\rho = 0$), the performance of our approach is even better than the performance of the next best approach (FixMatch) with $\rho = 1$ *i.e.*, when all unla-

Approach		10%	
Supervised (8f)		17.53 ± 0.49	
	$\rho = 1$	$\rho = 0.5$	$\rho = 0$
Pseudo-Label [35] (ICMLW' 13)	18.00 ± 0.16	$17.87 {\pm} 0.14$	17.79 ± 0.33
FixMatch [47] (NeurIPS' 20)	$18.02 {\pm} 0.31$	$18.00{\pm}0.29$	$17.96{\pm}0.25$
TCL (Ours)	19.13 ± 0.37	$18.95 {\pm} 0.17$	18.50 ± 0.95
TCL w/ Finetuning	$19.68 {\pm} 0.37$	$19.58{\pm}0.31$	$19.56{\pm}0.82$

Table 3: Semi-supervised action recognition under domain shift (Charades-Ego). Numbers show mean average precision (mAP) with ResNet-18 backbone across three different proportions of unlabeled data (ρ) between third and first person videos. **TCL** achieves the best mAP, even on this challenging dataset.



Figure 5: Qualitative examples comparing TCL with supervised baseline, S4L [60] and FixMatch [47]. Top Row: Top-1 predictions using ResNet-18 trained with 1% labeled data from Jester, Bottom Row: Top-1 predictions using ResNet-18 trained with 5% labeled data from Kinetics-400. TCL is able to correctly recognize different hand gestures in Jester and diverse human actions in Kinetics-400 dataset. (Best viewed in color.)

beled data is from the target domain itself. This depicts the robustness of **TCL** and its ability to harness diverse domain data more efficiently in semi-supervised setting.

Qualitative Results. Figure 5 shows qualitative comparison between our approach **TCL** and other competing methods (S4L [60] and FixMatch [47]) including the simple supervised baseline on Jester and Kinetics-400 respectively. As can be seen, our temporal contrastive learning approach is able to correctly recognize different hand gestures from Jester dataset even with 1% of labeling, while the supervised baseline and the next best approach (S4L) fail to recognize such actions. Similarly, our approach by effectively utilizing temporal information, predicts the correct label in most cases including challenging actions like 'climbing tree' and 'zumba' on Kinetics-400 dataset. More qualitative examples are included in the supplementary material.

Role of Pseudo-Labeling. We test the reliability of pseudolabeling on Jester (using ResNet-18 and 1% labeling) with 50 epoch intervals and observe that the pseudo-labeling accuracy gradually increases from 0% at the beginning to 65.95% at 100 epoch and then 93.23% at 350 epoch. This shows that while our model may create some wrong groups at the start, it gradually improves the groups as the training goes by, leading to a better representation by exploiting both instance and group contrastive losses.

Approach	Top-1 Accuracy
TCL w/o Group-Contrastive Loss	27.24 ± 0.42
TCL w/ Pseudo-Label Consistency Loss	$23.60{\pm}1.04$
TCL (Ours)	29.81±0.77

Table 4: Ablation Studies on Mini-Something-V2. Numbers show top-1 accuracy with ResNet-18 and 5% labeled Data.

4.3. Ablation Studies

We perform extensive ablation studies on Mini-Something-V2 with 5% labeled data and ResNet-18 backbone to better understand the effect of different losses and hyperparameters in our framework.

Effect of Group Contrastive Loss. We perform an experiment by removing group contrastive loss from our framework (ref. Section 3.2.3) and observe that top-1 accuracy drops to 27.24% from 29.81% (Table 4), showing the importance of it in capturing high-level semantics.

Ablation on Contrastive Loss. We investigate the effectiveness of our contrastive loss by replacing it with pseudolabel consistency loss used in FixMatch [47]. We observe that training with our contrastive loss, surpasses the performance of the training with the pseudo-label consistency loss by a high margin (around 6.21% gain in the top-1 accuracy) on Mini-Something-V2 (Table 4). We further compare our approach in the absence of group-consistency (TCL w/o Group-Contrastive Loss) with a variant of FixMatch [47] that uses temporal augmentation and observe that our approach still outperforms it by a margin of 2.66% (24.58%) vs 27.24%) on Mini-Something-V2 (with ResNet-18 and 5% labeling). This shows that temporal augmentation alone fails to obtain superior performance and this improvement is in fact due to the efficacy of our contrastive loss formulation over the pseudo-label loss used in FixMatch [47].

Effect of Different Frame Rate. We analyze the effect of doubling frame-rates in both pathways and observe that **TCL** (w/ 16 frame segments in base and 8 frame segments in the auxiliary pathway) improves top-1 accuracy by 1.5% on Mini-Something-V2 with ResNet-18 and 5% labeled data (29.81% vs 31.31%). However, due to heavy increase in compute and memory requirement with little relative gain in performance, we limit our study to 8 and 4 frame setting.

Effect of Hyperparameters. We analyze the effect of the ratio of unlabeled data to labeled data (μ) and observe that by setting μ to {3, 5, 7} with a fixed $\gamma = 1$, produces similar results on Mini-Something-V2 (Figure 6 (Left)). However, as scaling μ often requires high computational resources, we set it to 3 in all our experiments to balance the efficiency and accuracy in semi-supervised action recognition. We also find that weight of the instance-contrastive loss (γ) greatly affects the performance in semi-supervised learning as accuracy drops by more than 6% when setting γ to 3 in-



Figure 6: Effect of Hyperparameters on Mini-Something-V2. (Left) Varying the ratio of unlabeled data to the labeled data (μ), (Right) Varying the instance-contrastive loss weight (γ).

stead of the optimal value of 9 on Mini-Something-V2 with ResNet-18 backbone and 5% of labeling (Figure 6 (Right)).

Comparison With Self-Supervised Approaches. We compare our method with three video self-supervised methods, namely Odd-One-Out Networks (O3N) [18], Video Clip Order Prediction (COP) [57] and Memory-augmented Dense Predictive Coding (MemDPC) [26] through pretraining using self-supervised method and then finetuning using available labels on Mini-Something-V2 (with ResNet18 and 5% labeled data). Our approach significantly outperforms all the compared methods by a margin of 6%-10% (O3N: 19.56%, COP: 23.93%, MemDPC: 18.67%, TCL: 29.81%), showing its effectiveness over self-supervised methods. Moreover, we also replace our temporal contrastive learning with O3N and observe that accuracy drops to 24.58% from 29.81%, showing the efficacy of our contrastive learning formulation over the alternate video-based self-supervised method on Mini-Something-V2.

5. Conclusion

We present a novel temporal contrastive learning framework for semi-supervised action recognition by maximizing the similarity between encoded representations of the same unlabeled video at two different speeds as well as minimizing the similarity between different unlabeled videos run at different speeds. We employ contrastive loss between different video instances including groups of videos with similar actions to explore high-level action semantics within the neighborhood of different videos depicting different instances of the same action. We demonstrate the effectiveness of our approach on four standard benchmark datasets, significantly outperforming several competing methods.

Acknowledgements. This work was partially supported by the SERB Grant SRG/2019/001205. This work is also supported by the Intelligence Advanced Research Projects Activity (IARPA) via DOI/IBC contract number D17PC00341. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

References

- Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Pseudo-Labeling and Confirmation Bias in Deep Semi-Supervised Learning. In *International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2020. 2
- [2] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with Pseudo-Ensembles. In *Neural Information Processing Systems*, pages 3365–3373, 2014. 2
- [3] 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 *IEEE Conference on Computer Vision and Pattern Recognition*, pages 9922–9931, 2020. 3
- [4] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-Supervised Learning with Distribution Matching and Augmentation Anchoring. In *International Conference on Learning Representations*, 2019. 2
- [5] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A Holistic Approach to Semi-Supervised Learning. In *Neural Information Processing Systems*, pages 5050–5060, 2019. 2, 5, 6
- [6] Léon Bottou. Large-Scale Machine Learning with Stochastic Gradient Descent. In *COMPSTAT*, pages 177–186. Springer, 2010. 5
- [7] Joao Carreira and Andrew Zisserman. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 1, 2, 4
- [8] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-Supervised Learning (chapelle, o. et al., eds.; 2006)[Book Reviews]. *Transactions on Neural Networks*, 20(3):542–542, 2009. 2
- [9] Chun-Fu Chen, Rameswar Panda, Kandan Ramakrishnan, Rogerio Feris, John Cohn, Aude Oliva, and Quanfu Fan. Deep Analysis of CNN-based Spatio-temporal Representations for Action Recognition. ArXiv preprint ArXiv:2010.11757, 2020. 2, 5
- [10] Ting Chen, Simon Kornblith, M. Norouzi, and Geoffrey E. Hinton. A Simple Framework for Contrastive Learning of Visual Representations. *ArXiv preprint ArXiv:2002.05709*, 2020. 1, 2, 3, 5
- [11] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big Self-Supervised Models are Strong Semi-Supervised Learners. arXiv preprint arXiv:2006.10029, 2020. 3
- [12] Jinwoo Choi, Gaurav Sharma, Manmohan Chandraker, and Jia-Bin Huang. Unsupervised and Semi-Supervised Domain Adaptation for Action Recognition from Drones. In *IEEE Winter Conference on Applications of Computer Vision*, pages 1717–1726, 2020. 7
- [13] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative Unsupervised Feature Learning with Convolutional Neural Networks. In Neu-

ral Information Processing Systems, pages 766–774, 2014.

- [14] Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, and Pascal Vincent. Why Does Unsupervised Pre-training Help Deep Learning? *Journal of Machine Learning Research*, 11:625–660, 2010. 5
- [15] Quanfu Fan, Chun-Fu Richard Chen, Hilde Kuehne, Marco Pistoia, and David Cox. More Is Less: Learning Efficient Video Representations by Big-Little Network and Depthwise Temporal Aggregation. In *Neural Information Processing Systems*, pages 2261–2270, 2019. 1, 2
- [16] Christoph Feichtenhofer. X3D: Expanding Architectures for Efficient Video Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 203–213, 2020.
 1
- [17] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. SlowFast Networks for Video Recognition. In *IEEE International Conference on Computer Vision*, pages 6202–6211, 2019. 1, 2
- [18] Basura Fernando, Hakan Bilen, Efstratios Gavves, and Stephen Gould. Self-supervised video representation learning with odd-one-out networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3636–3645, 2017. 8
- [19] Gaurav Fotedar, Nima Tajbakhsh, Shilpa Ananth, and Xiaowei Ding. Extreme Consistency: Overcoming Annotation Scarcity and Domain Shifts. arXiv preprint arXiv:2004.11966, 2020. 3
- [20] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised Representation Learning by Predicting Image Rotations. arXiv preprint arXiv:1803.07728, 2018. 2
- [21] Daniel Gordon, Kiana Ehsani, Dieter Fox, and Ali Farhadi. Watching The World Go By: Representation Learning from Unlabeled Videos. arXiv preprint arXiv:2003.07990, 2020.
 3
- [22] Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and Benchmarking Self-Supervised Visual Representation Learning. In *IEEE International Conference* on Computer Vision, pages 6391–6400, 2019. 1
- [23] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, Florian Hoppe, Christian Thurau, Ingo Bax, and Roland Memisevic. The "Something Something" Video Database for Learning and Evaluating Visual Common Sense. In *IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 1, 5
- [24] Yves Grandvalet and Yoshua Bengio. Semi-Supervised Learning by Entropy Minimization. In *Neural Information Processing Systems*, pages 529–536, 2005. 2
- [25] Tengda Han, Weidi Xie, and Andrew Zisserman. Video Representation Learning by Dense Predictive Coding. In Proceedings of the IEEE International Conference on Computer Vision Workshops, pages 0–0, 2019. 3
- [26] Tengda Han, Weidi Xie, and Andrew Zisserman. Memoryaugmented Dense Predictive Coding for Video Representation Learning. In *European conference on computer vision*. Springer, 2020. 1, 8

- [27] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Learning Spatio-Temporal Features with 3D Residual Networks for Action recognition. In *IEEE International Conference* on Computer Vision Workshops, pages 3154–3160, 2017. 2, 4
- [28] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum Contrast for Unsupervised Visual Representation Learning. In *IEEE Conference on Computer Vi*sion and Pattern Recognition, pages 9729–9738, 2020. 3
- [29] Olivier J Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-Efficient Image Recognition with Contrastive Predictive Coding. arXiv preprint arXiv:1905.09272, 2019. 1
- [30] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning Deep Representations by Mutual Information Estimation and Maximization. arXiv preprint arXiv:1808.06670, 2018. 3
- [31] Alexandros Iosifidis, Anastasios Tefas, and Ioannis Pitas. Semi-supervised classification of human actions based on neural networks. In 2014 22nd International Conference on Pattern Recognition, pages 1336–1341. IEEE, 2014. 2
- [32] 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
- [33] Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative Learning of Audio and Video Models from Self-Supervised Synchronization. In *Neural Information Processing Systems*, pages 7763–7774, 2018. 3
- [34] Samuli Laine and Timo Aila. Temporal Ensembling for Semi-Supervised Learning. *arXiv preprint arXiv:1610.02242*, 2016. 2
- [35] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *International Conference on Machine Learning Work-shop*, volume 3, page 2, 2013. 1, 2, 5, 6, 7
- [36] Ji Lin, Chuang Gan, and Song Han. TSM: Temporal Shift Module for Efficient Video Understanding. In *IEEE International Conference on Computer Vision*, pages 7083–7093, 2019. 1, 2, 4, 5, 6
- [37] Joanna Materzynska, Guillaume Berger, Ingo Bax, and Roland Memisevic. The Jester Dataset: A Large-Scale Video Dataset of Human Gestures. In *IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. 1, 2, 5
- [38] Ishan Misra and Laurens van der Maaten. Self-Supervised Learning of Pretext-Invariant Representations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6707–6717, 2020. 1, 3
- [39] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual Adversarial Training: A Regularization Method for Supervised and Semi-Supervised Learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018. 2

- [40] Augustus Odena. Semi-Supervised Learning with Generative Adversarial Networks. arXiv preprint arXiv:1606.01583, 2016. 2
- [41] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal Contrastive Video Representation Learning. arXiv preprint arXiv:2008.03800, 2020. 1, 3
- [42] Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. In *Neural Information Processing Systems*, pages 3546–3554, 2015. 2
- [43] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, Sergey Levine, and Google Brain. Time-Contrastive Networks: Self-Supervised Learning from Video. In 2018 IEEE International Conference on Robotics and Automation (ICRA), pages 1134–1141. IEEE, 2018. 3
- [44] Gunnar A Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, and Karteek Alahari. Charades-Ego: A Large-Scale Dataset of Paired Third and First Person Videos. arXiv preprint arXiv:1804.09626, 2018. 2, 5, 7
- [45] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in Homes: Crowdsourcing Data Collection for Activity Understanding. In *European Conference on Computer Vision*, pages 510–526. Springer, 2016. 7
- [46] Karen Simonyan and Andrew Zisserman. Two-Stream Convolutional Networks for Action Recognition in Videos. In *Neural Information Processing Systems*, pages 568–576, 2014. 2
- [47] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. In *Neural Information Processing Systems*, 2020. 1, 2, 3, 5, 6, 7, 8
- [48] Antti Tarvainen and Harri Valpola. Mean Teachers are Better Role Models: Weight-Averaged Consistency Targets Improve Semi-Supervised Deep Learning Results. In *Neural Information Processing Systems*, pages 1195–1204, 2017. 2, 5, 6
- [49] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning Spatiotemporal Features with 3D Convolutional Networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489– 4497, 2015. 1, 2
- [50] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation Learning with Contrastive Predictive Coding. arXiv preprint arXiv:1807.03748, 2018. 1, 2, 3
- [51] Jiangliu Wang, Jianbo Jiao, and Yun-Hui Liu. Self-Supervised Video Representation Learning by Pace Prediction. arXiv preprint arXiv:2008.05861, 2020. 3
- [52] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal Segment Networks: Towards Good Practices for Deep Action Recognition. In *European conference on computer vision*, 2016. 2, 3, 4

- [53] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised Feature Learning via Non-Parametric Instance Discrimination. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 3733– 3742, 2018. 3
- [54] Fanyi Xiao, Yong Jae Lee, Kristen Grauman, Jitendra Malik, and Christoph Feichtenhofer. Audiovisual Slow-Fast Networks for Video Recognition. arXiv preprint arXiv:2001.08740, 2020. 1
- [55] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation for consistency training. arXiv preprint arXiv:1904.12848, 2019.
 2
- [56] Qizhe Xie, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Self-training with Noisy Student improves ImageNet classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 10687–10698, 2020. 5
- [57] Dejing Xu, Jun Xiao, Zhou Zhao, Jian Shao, Di Xie, and Yueting Zhuang. Self-Supervised Spatiotemporal Learning via Video Clip Order Prediction. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 10334– 10343, June 2019. 8
- [58] Ceyuan Yang, Yinghao Xu, Bo Dai, and Bolei Zhou. Video Representation Learning with Visual Tempo Consistency. arXiv preprint arXiv:2006.15489, 2020. 3, 5
- [59] Yuan Yao, Chang Liu, Dezhao Luo, Yu Zhou, and Qixiang Ye. Video Playback Rate Perception for Self-Supervised Spatio-Temporal Representation Learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 6548–6557, 2020. 3
- [60] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4L: Self-Supervised Semi-Supervised Learning. In *IEEE International Conference on Computer Vision*, pages 1476–1485, 2019. 2, 5, 6, 7
- [61] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal Relational Reasoning in Videos. In European Conference on Computer Vision (ECCV), pages 803– 818, 2018. 2
- [62] Barret Zoph, Golnaz Ghiasi, Tsung-Yi Lin, Yin Cui, Hanxiao Liu, Ekin D. Cubuk, and Quoc V. Le. Rethinking Pretraining and Self-training. In *Neural Information Processing Systems*, 2020. 5