

# TimeBalance: Temporally-Invariant and Temporally-Distinctive Video Representations for Semi-Supervised Action Recognition

Ishan Rajendrakumar Dave, Mamshad Nayeem Rizve, Chen Chen, Mubarak Shah  
Center for Research in Computer Vision, University of Central Florida, Orlando, USA  
{ishandave, nayeemrizve}@knights.ucf.edu, {chen.chen, shah}@csrcv.ucf.edu

## Abstract

*Semi-Supervised Learning can be more beneficial for the video domain compared to images because of its higher annotation cost and dimensionality. Besides, any video understanding task requires reasoning over both spatial and temporal dimensions. In order to learn both the static and motion related features for the semi-supervised action recognition task, existing methods rely on hard input inductive biases like using two-modalities (RGB and Optical-flow) or two-stream of different playback rates. Instead of utilizing unlabeled videos through diverse input streams, we rely on self-supervised video representations, particularly, we utilize temporally-invariant and temporally-distinctive representations. We observe that these representations complement each other depending on the nature of the action. Based on this observation, we propose a student-teacher semi-supervised learning framework, TimeBalance, where we distill the knowledge from a temporally-invariant and a temporally-distinctive teacher. Depending on the nature of the unlabeled video, we dynamically combine the knowledge of these two teachers based on a novel temporal similarity-based reweighting scheme. Our method achieves state-of-the-art performance on three action recognition benchmarks: UCF101, HMDB51, and Kinetics400. Code: <https://github.com/DAVEISHAN/TimeBalance>.*

## 1. Introduction

Recent development in action recognition have opened up a wide range of real-world applications: visual security systems [19, 50, 56], behavioral studies [35], sports analytics [46], elderly person fall detection systems [10, 49, 85], etc. Most of these developments are mainly courtesy of large-scale curated datasets like Kinetics [12], HVU [21], and HACS [87]. However, labeling such a massive video dataset requires an enormous amount of annotation time and human effort. At the same time, there is a vast

amount of unlabeled videos available on the internet. The goal of semi-supervised action recognition is to use such large-scale unlabeled dataset to provide additional supervision along with the labeled supervision of the small-scale dataset.

Semi-supervised learning for image classification has seen tremendous progress in recent years [1, 57, 63, 68]. In semi-supervised action recognition, recent approaches have adapted these image-based methods by incorporating motion-related inductive biases into the setup. For instance, some methods [77, 79] use two different input modalities where the original RGB video promotes learning appearance-based features while optical flow/temporal gradients promotes learning of motion-centric features. Another set of methods uses input streams of different sampling rates to achieve this [64, 70]. Although these input-level inductive biases are simple-yet-very-effective to provide unlabeled supervision for action recognition, they are not suitable for large-scale datasets due to their multiplicative storage requirement and high preprocessing overhead.

Contrastive Self-supervised Learning (CSL) has emerged as a powerful technique to learn meaningful representations from unlabeled videos. Existing video CSL methods deal with mainly two different kinds of objectives: (1) Learning similarities across clips of the same video i.e. *temporally-invariant* representations [29, 53, 55] (2) Learning differences across clips of the video i.e. *temporally-distinctive* representations [18, 36, 75]. Each objective has its own advantages, depending on the nature of the unlabeled videos being used.

Our experiments reveal a clear difference in the class-wise performance of both methods, as illustrated in Fig. 1. The right half of the figure shows the action classes where the temporally-invariant model is dominant. We can observe that all such action classes are atomic actions with high repetitions, e.g., Fencing, Knitting. Any two clips from such videos are highly similar, hence, increasing agreement between them i.e. learning *temporal-invariant* representation is more meaningful. The left half of the figure shows action classes where temporally-distinctive rep-

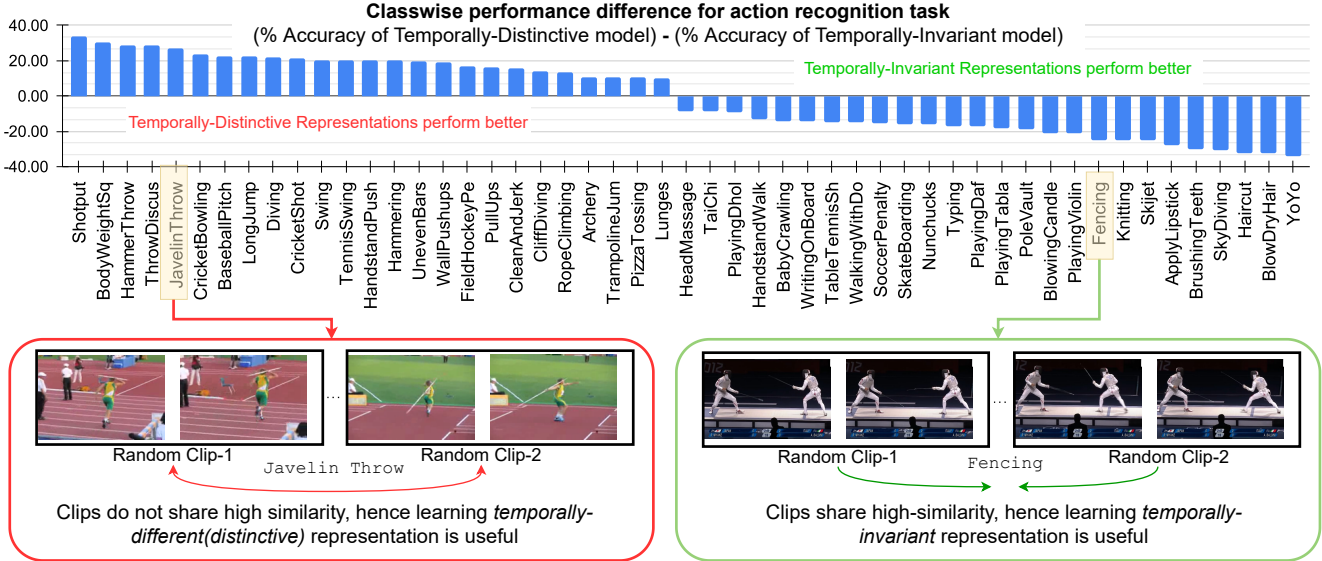


Figure 1. **Motivation for Temporally-Distinctive and Temporally-Invariant Representations.** In order to leverage the unlabeled videos effectively, we consider two kinds of self-supervised video representation learning techniques with complementary goals: (1) *Temporally Invariant Representations* (Bottom Right) encourage learning the commonalities of the clips, hence it mainly focuses on learning features related to highly frequent repetitions and appearance. (2) *Temporally Distinctive Representations* (Bottom Left) encourage learning the dissimilarities between clips of the same video, hence it encourages learning features for sub-actions within the video. The plot shows the activity-wise UCF101 performance difference of finetuned models which were self-supervised pretrained with temporally-distinctive and temporally-invariant objectives. The plot shows extreme 25-25 classes after sorting the all classwise differences.

representations perform better. We can observe that such action classes are slightly more complex i.e. they contain sub-actions, e.g., JavelinThrow first involves running and then throwing. Any two clips from such videos are visually very different, hence if we maximize agreement between them then it results in loss of the temporal dynamics. Therefore, *temporally-distinctive* representation is more suitable in such videos.

Based on our observation, we aim to leverage the strengths of both temporally-invariant and temporally-distinctive representations for semi-supervised action recognition. To achieve this, we propose a semi-supervised framework based on a student-teacher setup. The teacher supervision includes two models pre-trained using CSL with temporally-invariant and temporally-distinctive objectives. After pre-training, the teachers are fine-tuned with the labeled set to adapt to the semi-supervised training of the student. During semi-supervised training, we weigh each teacher model based on the nature of the unlabeled video instance. We determine the nature of the instance by computing its similarity score using the temporal self-similarity matrices of both teachers. This way, the student is trained using the labeled supervision from the labeled set and the unlabeled supervision from the weighted average of the teachers. It is worth noting that our framework doesn't depend on complicated data-augmentation schemes

like FixMatch [65].

The contributions of this work are summarized as follows:

- We propose a student-teacher-based semi-supervised learning framework that consists of two teachers with complementary self-supervised video representations: Temporally-Invariant and Temporally-Distinctive.
- In order to leverage the strength of each representation (invariant or distinctive), we weigh a suitable teacher according to an unlabeled video instance. We achieve it by the proposed temporal-similarity-based reweighting scheme.
- Our method outperforms existing approaches and achieves state-of-the-art results on popular action recognition benchmarks, including UCF101, HMDB51, and Kinetics400.

## 2. Prior Work

**Semi-supervised learning in images** Semi-supervised learning is one of the fundamental approaches to learning from limited labeled data [16, 26, 40, 41, 48, 54, 58]. There are two common approaches for semi-supervised learning are consistency regularization [43, 52, 61, 68, 78] and pseudo-labeling [1, 44, 57, 59, 63, 84]. The consistency regularization methods attempt to achieve perturbation invariant output space. To this end, these methods try to minimize a distance/divergence-based loss as a measure of

consistency between two differently perturbed/augmented versions of an image. Pseudo-labeling-based methods on the other hand promote entropy minimization to improve by performing self-training on the confident pseudo-labels generated from the network. Hybrid methods [7,8,65] combine consistency regularization and pseudo-labeling techniques to obtain further improvement. Some of the recent works [3, 16, 83] have also shown the effectiveness of self-supervised representations in solving this task. Our work is somewhat in the spirit of this last set of works as we also leverage self-supervised representation learning to obtain temporally distinct and invariant features.

**Semi-supervised learning in videos** Although there is a tremendous recent development in action recognition [2, 6, 9, 13, 20, 22, 32, 47, 51, 60, 72, 74, 86], semi-supervised learning in videos is not explored as in the image domain. CMPL [80] utilizes a FixMatch framework, where they study the effect of model capacity to provide complementary gains of unlabeled videos. They provide empirical evidence that the smaller model is responsible for learning the motion features whereas, the bigger model learns more appearance-biased features. However, defining smaller and bigger models is very relative, and the observation may not hold true for different architectures. Some prior work injects temporal dynamics to their semi-supervised framework by additional input modality to RGB videos like temporal gradients [77], optical-flow [79], or P-frames [69]. Another set of methods utilizes consistency loss between the slow and fast input streams of the video to leverage the unlabeled videos [64, 70]. We can see that the prior works rely on hard-inductive biases like model architectures, input modality, or input sampling to learn temporal-invariance and distinctiveness. On the other hand, we do not have such a hard design choice; we leverage the unlabeled videos by the nature of the video instance using two complementary self-supervised teachers.

**Self-supervised learning (SSL)** In recent years, self-supervised learning has demonstrated learning powerful representations for images [11, 15, 30, 33, 34, 82] and videos [18, 23, 25, 36, 37, 55, 62, 71]. Although some works [16, 83] have exploited self-supervised representation in semi-supervised image classification, Video SSL is not explored yet for semi-supervised action recognition.

Video SSL methods can be grouped mainly into two categories: the first set of methods focuses on learning temporal invariance [25, 53, 55]. These methods are a simple-yet-effective extension of instance-discrimination-based methods like SimCLR [15], MoCo [34], etc., where mutual information between two views of an instance is maximized. For videos, two views are two clips from different timestamps, hence it introduces temporal invariance in the learned representations. The second set of methods focuses on learning the temporal distinctiveness, where they try to

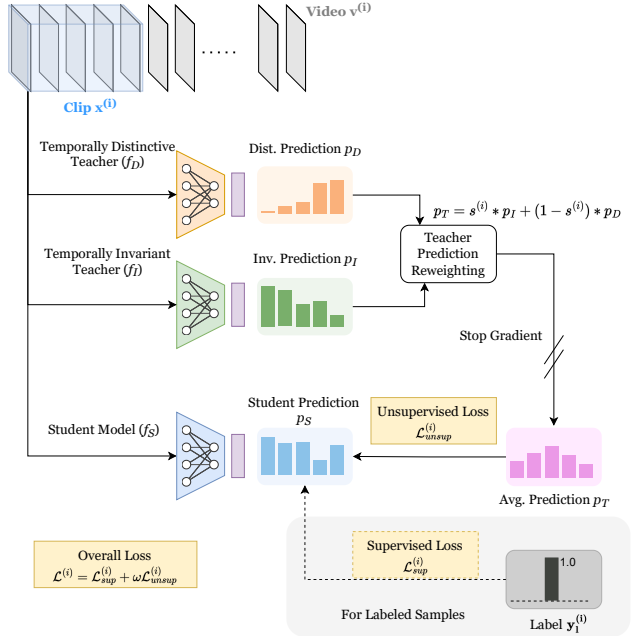


Figure 2. **Our Framework** We use a teacher-student framework where we use two teachers:  $f_I$  and  $f_D$ . The input clip  $\mathbf{x}^{(i)}$  is given to the teachers and student to get their predictions. We utilize a reweighting strategy to combine the predictions of two teachers. Regardless of whether the video  $\mathbf{v}^{(i)}$  is labeled or unlabeled, we distill the combined knowledge of teachers to the student. For the labeled samples, we also apply standard cross-entropy loss.

learn different representations for different clips through contrastive loss [18, 67, 75, 81] or through different temporal pretext transformations [5, 14, 17, 36, 37, 76].

TCLR [18] introduces temporal contrastive losses for both temporally pooled and unpooled features to learn the temporal distinctiveness. In our method, we utilized these losses to learn the temporal distinctive teacher.

### 3. Method

Let’s consider a small labeled set of videos  $\mathbb{D}_l = \{(\mathbf{v}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^{N_l}$ , where  $\mathbf{v}^{(i)}$  and  $\mathbf{y}^{(i)}$  denote  $i$ th video instance and its associated action label and  $N_l$  is number of total instances in the dataset. We also have access to a unlabeled dataset  $\mathbb{D}_u = \{\mathbf{v}^{(i)}\}_{i=1}^{N_u}$ , where  $N_u$  is the total number of unlabeled videos and  $N_u \gg N_l$ . The objective of the semi-supervised action recognition is to utilize both labeled and unlabeled sets ( $\mathbb{D}_l$  and  $\mathbb{D}_u$ ) to improve the action recognition performance.

A high-level schematic diagram of our framework is depicted in Fig. 2. Our semi-supervised learning framework, *TimeBalance*, is a teacher-student framework. To train on the unlabeled samples, we distill the knowledge of two teacher models: temporally-invariant teacher  $f_I$  and temporally-distinctive teacher  $f_D$ , which are trained in self-

supervised manner, to a student model  $f_S$ . In the following, we explain the details of *TimeBalance*. To be particular, in Sec. 3.1, we present the self-supervised pretraining of teacher models. In Sec. 3.2 we explain the semi-supervised training of the student model and our Temporal Similarity based Teacher Reweighting (TSTR) scheme.

### 3.1. Self-supervised pretraining of teachers

From a video instance  $\mathbf{v}^{(i)}$ , we sample  $n$  consecutive clips  $\mathbb{X}^{(i)} = \{\mathbf{x}_t^{(i)}\}_{t=1}^n$ , where  $t$  represents clip location (timestamp). Each of these clips undergoes stochastic transformations (e.g. random crop, random color jittering, etc.). Next, we send these clips to the teacher model  $f$  and a non-linear projection head  $g$  respectively. The non-linear projection head, projects the clip-embedding from the teacher model to a lower-dimensional normalized representation  $\mathbf{z}$ , s.t.  $\mathbf{z} \in \mathbb{R}^d$ , where  $d$  is the dimension of output vector  $\mathbf{z}$ .

#### 3.1.1 Pretraining of Temporally-Invariant Teacher

The goal of temporal-invariant pretraining is to learn the shared information across the  $n$  different clips  $\{\mathbf{x}_t^{(i)}\}_{t=1}^n$  of the same video instance  $i$ . To achieve this, we maximize the agreement between the projections,  $\mathbf{z}$ , of two different clips from the same video instance  $\{(\mathbf{z}_{t_1}^{(i)}, \mathbf{z}_{t_2}^{(i)}) \mid t_1, t_2 \in \{1 \dots n\} \text{ and } t_1 \neq t_2\}$ , while maximizing the disagreement between the projections of clips from different video instances  $\{(\mathbf{z}^{(i)}, \mathbf{z}^{(j)}) \mid i, j \in B \text{ and } i \neq j\}$ , where  $B$  is the batch-size. This contrastive objective can be expressed as the following equation:

$$\mathcal{L}_I^{(i)} = - \sum_{\substack{t_1, t_2 \\ t_2 \neq t_1}}^n \log \frac{h(\mathbf{z}_{t_1}^{(i)}, \mathbf{z}_{t_2}^{(i)})}{\sum_{j=1}^B \mathbb{1}_{[j \neq i]} h(\mathbf{z}_{t_1}^{(i)}, \mathbf{z}_{t_1}^{(j)}) + h(\mathbf{z}_{t_1}^{(i)}, \mathbf{z}_{t_2}^{(j)})}, \quad (1)$$

where  $h(\mathbf{u}_1, \mathbf{u}_2) = \exp(\mathbf{u}_1^T \mathbf{u}_2 / (\|\mathbf{u}_1\| \|\mathbf{u}_2\| \tau))$  is used to compute the similarity between  $\mathbf{u}_1$  and  $\mathbf{u}_2$  vectors with an adjustable temperature parameter,  $\tau$ .  $\mathbb{1}_{[j \neq i]} \in \{0, 1\}$  is an indicator function which equals 1 iff  $j \neq i$ .

$$\mathcal{L}^{(i)} = \mathcal{L}_{sup}^{(i)} + \omega \mathcal{L}_{unsup}^{(i)}, \quad (2)$$

where  $\omega$  is the weight of the unsupervised loss.

#### 3.1.2 Pretraining of Temporally-Distinctive Teacher

Contrary to the goal of  $\mathcal{L}_I$ , temporally-distinctive pretraining deals with learning the differences across the clips of the same video instance. To achieve this, we generate another set of clips  $\tilde{\mathbb{X}}^{(i)} = \{\tilde{\mathbf{x}}_t^{(i)}\}_{t=1}^n$ , which are randomly-augmented versions of clips in  $\mathbb{X}^{(i)}$ . After that, we maximize the agreement between the projections of a pair of clips  $\{(\mathbf{z}_{t_1}^{(i)}, \tilde{\mathbf{z}}_{t_1}^{(i)}) \mid t_1 \in \{1 \dots n\}\}$  from the same timestamp

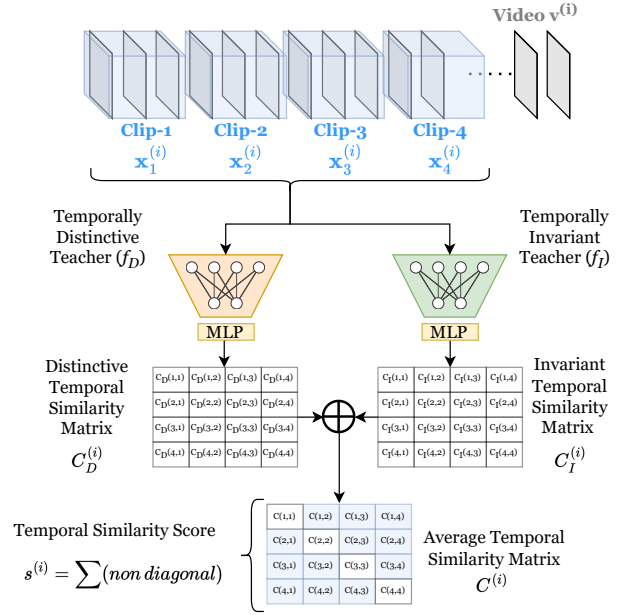


Figure 3. **Temporal Similarity based Teacher Reweighting** Firstly, a set of clips from video  $\mathbf{v}^{(i)}$  are passed through the distinctive and invariant teachers to get representations. Secondly, Temporal Similarity Matrices ( $C_D^{(i)}$  and  $C_I^{(i)}$ ) are constructed from the cosine similarity of one timestamp to another timestamp. Finally, from the average matrix  $C^{(i)}$ , a temporal similarity score  $s^{(i)}$  is computed.  $s^{(i)}$  is utilized to combine predictions of teachers for video  $\mathbf{v}^{(i)}$  during semi-supervised training.

and maximize the disagreement between the projections of pair of temporally misaligned clips. A mathematical expression for this contrastive objective can be written as:

$$\mathcal{L}_{D1}^{(i)} = - \sum_{t_1=1}^n \log \frac{h(\mathbf{z}_{t_1}^{(i)}, \tilde{\mathbf{z}}_{t_1}^{(i)})}{\sum_{\substack{t_2=1 \\ t_2 \neq t_1}}^n h(\mathbf{z}_{t_1}^{(i)}, \mathbf{z}_{t_2}^{(i)}) + h(\mathbf{z}_{t_1}^{(i)}, \tilde{\mathbf{z}}_{t_2}^{(i)})}, \quad (3)$$

The above contrastive loss imposes temporal distinctiveness at the clip-level i.e. temporally-average-pooled features. Similarly, we can also impose temporal distinctiveness ( $\mathcal{L}_{D2}^{(i)}$ ) on a more fine-grained level i.e. on the unpooled temporal feature slices [18]. More details in Supp. Sec. D. We combine these pooled and unpooled temporal-distinctiveness objectives to obtain  $\mathcal{L}_D = \mathcal{L}_{D1} + \mathcal{L}_{D2}$ .

The primary objective of this work is semi-supervised action recognition. Therefore, even though the self-supervised teacher models lack any explicit notion of category-specific output space, we argue that to solve the downstream action recognition task their knowledge has to be distilled from the action category-specific output space. To this end, we finetune both of the self-supervised teacher



models on the labeled set  $\mathbb{D}_l$  using cross-entropy loss.

### 3.2. Semi-supervised training of student model

We initialize the student model with weights from a video self-supervised model [18] trained on  $\mathbb{D}_u$ . During the semi-supervised training, student model  $f_S$  gets supervision from two sources: (i) ground-truth label (if available), and (ii) teacher supervision (Fig. 2). A visual aid is provided in [Supp. Sec. D](#) for loss computations over the labeled and unlabeled split in semi-supervised training. The supervision from the ground-truth label is utilized using standard cross-entropy loss, as depicted in the following equation:

$$\mathcal{L}_{sup}^{(i)} = - \sum_{c=1}^C \mathbf{y}_c^{(i)} \log \mathbf{p}_c^{(i)} \quad (4)$$

For instance  $i$ , the prediction vectors of the invariant and distinctive teacher are denoted as  $\mathbf{p}_I$  and  $\mathbf{p}_D$ , respectively.

Next, we will discuss our TSTR scheme to distill knowledge from the temporally invariant and distinct teachers to train the student model on the unlabeled data,  $\mathbb{D}_u$  and videos of labeled data  $\mathbb{D}_l$ .

**Temporal Similarity based Teacher Reweighting.** In order to combine supervision from  $f_I$  and  $f_D$  for a particular video instance  $\mathbf{v}^{(i)}$ , we first compute temporal similarity scores. To this end, we compute the cosine similarity between each pair of clips to form a temporal similarity matrix,  $\mathbf{C}^{(i)}$ , as depicted in Fig. 3. The temporal similarity matrix computation is described in the following,

$$\mathbf{C}^{(i)} = \left[ \text{Sim}(\mathbf{z}_{t_1}^{(i)}, \mathbf{z}_{t_2}^{(i)}) \right]_{t_1, t_2=1}^n, \quad (5)$$

where,  $\text{Sim}(\cdot)$  is the cosine similarity function.

We compute the temporal similarity matrix for both invariant and distinctive teachers denoted as  $\mathbf{C}_I^{(i)}$ , and  $\mathbf{C}_D^{(i)}$ , respectively. Next, in order to get a similarity score,  $s^{(i)}$ , for an instance  $i$ , we take the average of non-diagonal elements of both  $\mathbf{C}_I^{(i)}$ , and  $\mathbf{C}_D^{(i)}$  matrices.

$$s^{(i)} = \frac{1}{2n(n-1)} \sum_{\substack{t_1, t_2=1 \\ t_2 \neq t_1}}^n (\mathbf{C}_I^{(i)} + \mathbf{C}_D^{(i)}) \quad (6)$$

We use this temporal similarity score,  $s^{(i)}$ , to aggregate the outputs of the teacher models. Let's assume, for instance  $i$ , the prediction vectors of the invariant and distinctive teacher are denoted as  $\mathbf{p}_I^{(i)}$  and  $\mathbf{p}_D^{(i)}$ , respectively. Now, we want to combine these teacher prediction vectors in such a way that the temporally-invariant prediction  $\mathbf{p}_I^{(i)}$  gets a higher weight if the temporal similarity score is high, and the weight of the temporal-distinctive prediction  $\mathbf{p}_D^{(i)}$  gets higher in the case of a lower temporal similarity score.

This dynamic weighting scheme for obtaining the aggregated teacher prediction is provided below.

$$\mathbf{p}_T^{(i)} = s^{(i)} \cdot \mathbf{p}_I^{(i)} + (1 - s^{(i)}) \cdot \mathbf{p}_D^{(i)} \quad (7)$$

We use the combined teacher prediction  $\mathbf{p}_T^{(i)}$  to provide supervision to the student model using  $\mathcal{L}_2$  loss. We compute this loss only on the unlabeled samples ( $\mathbb{D}_u$  and videos of  $\mathbb{D}_l$  without assigned labels), hence, we refer to this loss as the unsupervised loss. This loss term is defined below.

$$\mathcal{L}_{unsup}^{(i)} = \left( \mathbf{p}_T^{(i)} - \mathbf{p}_S^{(i)} \right)^2 \quad (8)$$

Finally, we sum both the supervised and unsupervised losses to train the student model. The overall objective function is defined below.

### 3.3. Algorithm

Let's consider models  $f_I$ ,  $f_D$ , and  $f_S$  are parameterized by  $\theta_I$ ,  $\theta_D$ , and  $\theta_S$ . All steps of our semi-supervised training are put together in Algorithm 1.

---

#### Algorithm 1: TimeBalance training algorithm

---

- 1 **Inputs:**
  - 2 *Datasets:*  $\mathbb{D}_u, \mathbb{D}_l$
  - 3 *#Epochs:*  $max\_ssl\_epoch_I, max\_ssl\_epoch_D, max\_epoch_{tune}, max\_epoch$
  - 4 *Learning Rates:*  $\alpha_I, \alpha_D, \alpha_S$
  - 5 **Output:** Student model  $\theta_S$
  - 6 Initialize  $\theta_I, \theta_D$  randomly;
  - 7 Initialize  $\theta_S$  with any SSL [18, 55] method on  $\mathbb{D}_u$ ;
  - 8 Temporally-Invariant Self-supervised Pretraining:
  - 9 **for**  $e_0 \leftarrow 1$  **to**  $max\_ssl\_epoch_I$  **do**
  - 10 |  $\theta_I \leftarrow \theta_I - \alpha_I \nabla_{\theta_I} L_I(\theta_I)$
  - 11 **end**
  - 12 Temporally-Distinctive Self-supervised Pretraining:
  - 13 **for**  $e_0 \leftarrow 1$  **to**  $max\_ssl\_epoch_D$  **do**
  - 14 |  $\theta_D \leftarrow \theta_D - \alpha_D \nabla_{\theta_D} L_D(\theta_D)$
  - 15 **end**
  - 16 Compute the similarity score  $s^{(i)}$  using Eq. 6
  - 17 Finetuning the teacher models on labeled set  $\mathbb{D}_l$ :
  - 18 **for**  $e_0 \leftarrow 1$  **to**  $max\_epoch_{tune}$  **do**
  - 19 |  $\theta_I^* = \text{argmin}_{\theta_I} L_{sup}(\theta_I)$
  - 20 |  $\theta_D^* = \text{argmin}_{\theta_D} L_{sup}(\theta_D)$
  - 21 **end**
  - 22 Semi-supervised training of student on  $\mathbb{D}_l + \mathbb{D}_u$ :
  - 23 **for**  $e_0 \leftarrow 1$  **to**  $max\_epoch$  **do**
  - 24 |  $L = L_{sup}(\theta_S) + \omega L_{unsup}(\theta_S, \theta_D^*, \theta_I^*, s^{(i)})$
  - 25 |  $\theta_S \leftarrow \theta_S - \alpha_S \nabla_{\theta_S} L$
  - 26 **end**
-

Method	Backbone	Params (M)	Input	#F	UCF101					HMDB51			Kinetics400	
					1%	5%	10%	20%	50%	40%	50%	60%	1%	10%
PL <small>ICML'13</small> [45]	3D-ResNet18	13.5	V	16	-	17.6	24.7	37.0	47.5	27.3	32.4	33.5	-	-
MT <small>NeuRIPS'17</small> [68]	3D-ResNet18	13.5	V	16	-	17.5	25.6	36.3	45.8	27.2	30.4	32.2	-	-
S4L <small>ICCV'19</small> [83]	3D-ResNet18	13.5	V	16	-	22.7	29.1	37.7	47.9	29.8	31.0	35.6	-	-
UPS <small>ICLR'21</small> [57]	3D-ResNet18	13.5	V	16	-	-	-	39.4	50.2	-	-	-	-	-
SD <small>ICCV'19</small> [28]	3D-ResNet18	13.5	V	16	-	31.2	40.7	45.4	53.9	32.6	35.1	36.3	-	-
MT+SD <small>WACV'21</small> [38]	3D-ResNet18	13.5	V	16	-	30.3	40.5	45.5	53.0	32.3	33.6	35.7	-	-
3DRotNet <small>Arxiv'19</small> [39]	3D-ResNet18	13.5	V	16	15.0	31.5	40.4	47.1	-	-	-	-	-	-
VideoSemi <small>WACV'21</small> [38]	3D-ResNet18	13.5	V	16	-	32.4	42.0	48.7	54.3	32.7	36.2	37.0	-	-
TCL <small>CVPR'21</small> [64]	TSM-ResNet18	-	V	8	-	-	-	-	-	-	-	-	11.6	-
TG-FixMatch <small>CVPR'21</small> [77]	3D-ResNet18	13.5	V	8	-	<u>44.8</u>	62.4	<u>76.1</u>	<u>79.3</u>	<u>46.5</u>	<u>48.4</u>	<u>49.7</u>	9.8	43.8
MvPL <small>ICCV'21</small> [79]	3D-ResNet18	13.5	VFG	8	-	41.2	55.5	64.7	65.6	30.5	33.9	35.8	5.0	36.9
TCLR <small>CVIU'22</small> [18]	3D-ResNet18	13.5	V	16	<u>26.9</u>	-	66.1	73.4	76.7	-	-	-	-	-
CMPL <small>CVPR'22</small> [80]	3D-ResNet18	13.5	V	8	23.8	-	<u>67.6</u>	-	-	-	-	-	<u>16.5</u>	<u>53.7</u>
TACL <small>TSVT'22</small> [70]	3D-ResNet18	13.5	V	16	-	35.6	50.9	56.1	65.8	34.6	37.2	39.5	-	-
TACL <small>TSVT'22</small> [70]	3D-ResNet18	13.5	V	16	-	43.7	55.6	59.2	67.2	38.7	40.2	41.7	-	-
MemDPC <small>ECCV'20</small> [31]	3D-ResNet18	13.5	V	16	-	-	44.2	50.9	62.3	-	-	-	-	-
MotionFit <small>ICCV'21</small> [27]	3D-ResNet18	13.5	VF	16	-	-	-	57.7	59.0	-	-	-	-	-
Ours ( <i>TimeBalance</i> )	3D-ResNet18	13.5	V	8	<b>29.1</b>	<b>47.9</b>	<b>69.8</b>	<b>79.1</b>	<b>83.3</b>	<b>49.8</b>	<b>51.4</b>	<b>53.1</b>	<b>17.1</b>	<b>54.9</b>
ActorCM <small>Arxiv'21</small> [88]	R(2+1)D-34	33.3	V	8	-	<u>27.0</u>	40.2	<u>51.7</u>	<u>59.9</u>	<u>32.9</u>	<u>38.2</u>	<u>38.9</u>	-	-
ActorCM <small>Arxiv'21</small> [88]	R(2+1)D-34	33.3	V	8	-	45.1	53.0	57.4	64.7	35.7	39.5	40.8	-	-
FixMatch <small>NeuRIPS'20</small> [65]	SlowFast-R50	60	V	8	16.1	-	55.1	-	-	-	-	-	10.1	49.4
MvPL <small>ICCV'21</small> [79]	3D-ResNet50	31.8	VFG	8	22.8	-	<u>80.5</u>	-	-	-	-	-	17.0	58.2
CMPL <small>CVPR'22</small> [80]	3D-ResNet50	31.8	V	8	<u>25.1</u>	-	79.1	-	-	-	-	-	<u>17.6</u>	<u>58.4</u>
Ours ( <i>TimeBalance</i> )	3D-ResNet50	31.8	V	8	<b>30.1</b>	<b>53.5</b>	<b>81.1</b>	<b>83.3</b>	<b>85.0</b>	<b>52.6</b>	<b>53.9</b>	<b>54.5</b>	<b>19.6</b>	<b>61.2</b>

Table 1. **Comparison with state-of-the-art methods:** Methods using pre-trained ImageNet weights are shown in Grey. V- Video (RGB), F- Optical Flow, G- Temporal Gradients. Best results are shown in **Red**, and Second-best in **Blue**

## 4. Experiments

### 4.1. Datasets

**UCF101** [66] is a dataset for human action recognition collected from internet videos consisting 101 action classes. We use split-1 for our experiments, which has 9,537 train videos and 3,783 test videos.

**HMDB51** [42] is relatively a smaller dataset collected from movie videos. It has 51 human activity classes and has a high intra-class variance. We use split-1 in this paper, which has 3,570 train videos and 1,530 test videos.

**Kinetics400** [12] is a large-scale dataset collected from YouTube videos. It has a standard split of 240k training videos and 20k validation videos which covers 400 actions.

### 4.2. Implementation Details

For our default experimental setup, we use an input clip resolution of  $224 \times 224$ . We use 3D-ResNet-50 [24] for both student and teacher models. We use most commonly used augmentations: geometry-based augmentations like random cropping, and flipping, and color-based augmentations like random greyscale, and color jittering.

**Self-supervised pretraining** The pretraining is performed with clips of 16 frames for 100 epochs for Kinetics400, 250

epochs for UCF101, and HMDB51 experiments. For contrastive losses, the default temperature is set to 0.1.

**Semi-supervised training** Semi-supervised training is performed with clips of 8 frames for 150 epochs.

**Inference** We follow standard protocol [73] of averaging the predictions of 10 uniformly spaced clips and 3 different spatial scales to get a video-level prediction.

More implementation details in [Supp. Sec. B](#).

### 4.3. Comparison with prior work

We compare our method with image-based baselines, video self-supervised baselines, and recent state-of-the-art methods for semi-supervised action recognition in Table 1. We present the results in two sections based on the backbone architecture used: (1) 3D-ResNet18 and (2) 3D-ResNet50. We report Top-1 classification accuracy as the performance measure and follow the standard protocol of reporting average performance over three independent runs.

**Image-based baselines** We consider widely used semi-supervised image classification baselines like Psuedo-Label [45], MeanTeacher [68], UPS [57], S4L [83] and FixMatch [65]. From Table 1, we observe that the results of these image-based methods are significantly lower than the video-based methods across all benchmarks; which suggests that spatial information is not enough to excel in semi-

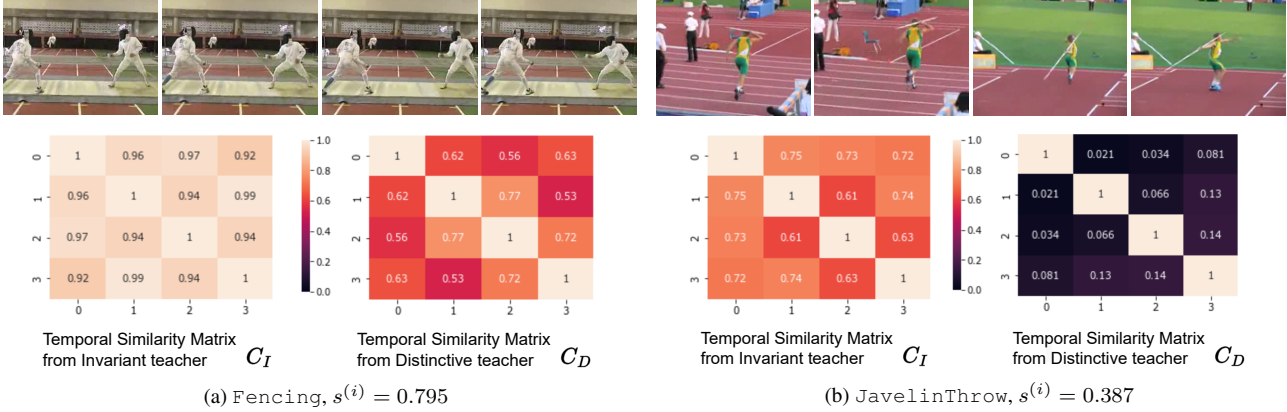


Figure 4. Visualization of similarity matrices from temporally-invariant and distinctive teachers and resultant instance-similarity score. (a) Video instances with atomic actions like `Fencing` result in a high instance similarity score, which will proportionally increase the weightage of temporally-invariant teacher (b) Whereas, complex actions video instances like `JavelinThrow` results in a low similarity score, which results into more weightage of temporally-distinctive teacher. Details of the computation of the similarity score in Sec 3.2

supervised action recognition.

**Video Self-supervised Learning baselines** Video self-supervised learning methods are first pretrained on the full-training set without using any labels and then finetuned on the labeled set in a supervised manner. We compare with contrastive learning-based methods like TCLR [18], MemDPC [31], multimodal approach like MotionFit [27], and pretext-task based approach 3DRotNet [39]. From Table 1, we notice that the performance of these video self-supervised methods is significantly better than image-based baselines, and in some cases, it even performs favorably against the semi-supervised action recognition methods [77, 79, 80]. However, our proposed semi-supervised method outperforms all these video self-supervised baselines by a noticeable margin.

**Semi-supervised action recognition baselines** We also compare with prior semi-supervised action recognition works like TG-Fixmatch [77], VideoSemi [38], MvPL [79], TACL [70], and ActorCM [89]. Our method achieves significant improvement over these methods across all benchmarks. Remarkably, our method also outperforms the methods that use additional modalities like optical flow in MvPL [79] and additional data (ImageNet) in ActorCM [89] and TACL [70].

#### 4.4. Ablations and Analysis

In the default setting of ablations, we consider UCF101 with 3D-ResNet50 as the teacher backbone. More ablations in Supp. Sec. C.

**Contribution of different training components** In Table 2, we analyze the effect of each model ( $f_S$ ,  $f_I$ ,  $f_D$ ) and our proposed teacher reweighting scheme with 5% and 20% labeled data on the UCF101 dataset. Row a-c, demonstrates the performance of teacher and student models indi-

vidually. Results in Row f demonstrates that the student model performs the best when we average the predictions of both teachers and removing any of them (Row d, e) degrades the performance. This validates our hypothesis that for optimal performance, we need to distill the knowledge from both teachers. Finally, Row g, demonstrates the effectiveness of our proposed teacher reweighting scheme using temporal similarity. We found similar results with  $f_S$  from random initialization in Table 3.

	$f_S$	$f_I$	$f_D$	Teacher Reweighting	UCF101 % labels
					5%    20%
(a)	✓	✗	✗	✗	48.66    79.70
(b)	✗	✓	✗	✗	43.79    74.55
(c)	✗	✗	✓	✗	44.08    75.13
(d)	✓	✓	✗	✗	47.95    79.28
(e)	✓	✗	✓	✗	48.81    79.76
(f)	✓	✓	✓	✗	52.14    82.02
(g)	✓	✓	✓	✓	53.48    83.26

Table 2. Ablation of different components of our framework. Student ( $f_S$ ) and teachers ( $f_I$  and  $f_D$ ) are 3D-ResNet50.

**Visualization of Temporal Similarity Matrix** We sample four consecutive frames from unlabeled videos and compute the temporal similarity matrix for both  $f_I$  and  $f_D$ . Visualization is shown Fig. 4. Firstly, we observe that  $f_I$  gives higher similarity values in both video instances than  $f_D$ . Secondly, we can observe that the final instance similarity score aligns with our goal i.e. providing more weight to  $f_I$  in atomic actions and providing more weight to  $f_D$  in complex actions.

**Different types of teacher pairs** In order to further investigate the importance of having teachers with complemen-

	$f_S$		Teacher		UCF101 % Labels	
	(rand. init.)	$f_I$ $f_D$	Reweighting	5%	20%	
(a)	✓	✗ ✗	✗	25.60	46.20	
(b)	✓	✓ ✗	✗	43.94	74.85	
(c)	✓	✗ ✓	✗	44.30	75.22	
(d)	✓	✓ ✓	✗	49.57	80.06	
(e)	✓	✓ ✓	✓	53.10	83.00	

Table 3. Ablation of different components with Student with random initialization. All backbones are 3D-ResNet50.

tary self-supervised objectives (i.e temporally invariant and distinctive), we study the effect of having various pairs of teachers in our framework (Table 4). *Inv1* and *Inv2* models are obtained from independent SSL pretrainings with temporal-invariance objective. Similarly, *Dist1* and *Dist2* are trained with temporal-distinctiveness SSL objective. We observe that, compared to combining teachers with the same SSL objective (Row a, b), teachers with different SSL objectives (Row c, e) perform significantly better. In our default setting, we use (*Inv1*, *Dist1*) as teacher pair.

	Teacher		UCF101 % labels	
	Teacher-1	Teacher-2	5%	20%
(a)	Inv1	Inv2	48.33	78.76
(b)	Dist1	Dist2	49.15	80.49
(c)	Inv1	Dist1	52.14	82.02
(d)	Inv1	Dist2	51.78	81.43

Table 4. Ablation of different teacher combinations.

**Student initialization** We also study various video self-supervised methods to initialize the student model and report the results in Fig 5. We conduct these experiments on the UCF101 dataset with 3D-ResNet18 as the student backbone. We observe that even across a diverse set of video SSL-based initialization techniques, our proposed semi-supervised framework can achieve better/competitive performance against the current state-of-the-art. This further validates the effectiveness of our idea of leveraging temporally distinct and invariant teachers and our proposed teacher prediction reweighting scheme.

**Number of clips in Teacher SSL pretraining** We analyze the effect of using a different number of clips to impose temporal invariance/distinctiveness in teacher pretraining in Table 5. We perform pretraining on UCF101 with the 3D-ResNet50 model and report the performance of the teacher model after finetuning on the labeled set. In the default setting, we use n=4 clips. We observe that while reducing the clips from 4 to 2, the performance drop is significant for temporally-distinctive teachers, whereas, it is not severe for the temporally-invariant teacher. Since the distinctiveness




Figure 5. Analysis of different student initialization strategies on UCF101 dataset with 5% labeled data and 3D-ResNet18 student.

contrastive loss  $\mathcal{L}_D$  treats other clips as negatives, its contrastive objective becomes very easy to solve if we reduce the number of clips (Eq. 3). Whereas, temporal-invariance contrastive loss  $\mathcal{L}_I$  takes negatives from the other instances of the batch, and decreasing the number of clips does not change the difficulty of the loss.

Pretraining	Number of clips (n)	UCF101 % Labels	
		5%	20%
Invariant Teacher	n=4	43.79	74.55
	n=2	42.78	73.86
Distinctive Teacher	n=4	44.08	75.13
	n=2	39.55	71.21

Table 5. Number of clips per video in self-supervised pretraining.

## 5. Conclusion and Future Work

In this work, we have proposed TimeBalance, a teacher-student framework for semi-supervised action recognition. We utilize the complementary strengths of temporally-invariant and temporally-distinctive representations to leverage unlabeled videos. Our extensive experimentation has empirically validated the effectiveness of different components of our framework and results on multiple benchmark datasets established TimeBalance as the new state-of-the-art for semi-supervised action recognition.

Our findings regarding the complementary strengths of temporally-invariant and temporally-distinctive video representations could be applied to other data-efficient video understanding problems, such as few-shot action recognition and spatio-temporal action detection. It would also be interesting to explore the invariance and distinctiveness properties of video using recent masked-image modeling techniques [23, 71] and multi-modal settings [4], such as audio + video and text + video.

## Acknowledgments

We thank Tristan de Blegiers for his help with visualization.



## References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2020. 1, 2
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6836–6846, 2021. 3
- [3] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Armand Joulin, Nicolas Ballas, and Michael Rabbat. Semi-supervised learning of visual features by non-parametrically predicting view assignments with support samples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8443–8452, 2021. 3
- [4] Roman Bachmann, David Mizrahi, Andrei Atanov, and Amir Zamir. MultiMAE: Multi-modal multi-task masked autoencoders. 2022. 8
- [5] Sagie Benaim, Ariel Ephrat, Oran Lang, Inbar Mosseri, William T Freeman, Michael Rubinstein, Michal Irani, and Tali Dekel. Speednet: Learning the speediness in videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9922–9931, 2020. 3
- [6] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, volume 2, page 4, 2021. 3
- [7] David Berthelot, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remix-match: Semi-supervised learning with distribution matching and augmentation anchoring. In *International Conference on Learning Representations*, 2020. 3
- [8] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems 32*, pages 5049–5059. Curran Associates, Inc., 2019. 3
- [9] Adrian Bulat, Juan Manuel Perez Rua, Swathikiran Sudhakaran, Brais Martinez, and Georgios Tzimiropoulos. Space-time mixing attention for video transformer. *Advances in Neural Information Processing Systems*, 34:19594–19607, 2021. 3
- [10] Marco Buzzelli, Alessio Albé, and Gianluigi Ciocca. A vision-based system for monitoring elderly people at home. *Applied Sciences*, 10(1):374, 2020. 1
- [11] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 9912–9924. Curran Associates, Inc., 2020. 3
- [12] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 1, 6
- [13] Jiawei Chen and Chiu Man Ho. Mm-vit: Multi-modal video transformer for compressed video action recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1910–1921, 2022. 3
- [14] Peihao Chen, Deng Huang, Dongliang He, Xiang Long, Runhao Zeng, Shilei Wen, Mingkui Tan, and Chuang Gan. Rspnet: Relative speed perception for unsupervised video representation learning. In *The AAAI Conference on Artificial Intelligence (AAAI)*, 2021. 3
- [15] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 3
- [16] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. *Advances in Neural Information Processing Systems*, 33, 2020. 2, 3
- [17] Hyeon Cho, Taehoon Kim, Hyung Jin Chang, and Wonjun Hwang. Self-supervised visual learning by variable playback speeds prediction of a video. *IEEE Access*, 9:79562–79571, 2021. 3
- [18] Ishan Dave, Rohit Gupta, Mamshad Nayeem Rizve, and Mubarak Shah. Tclr: Temporal contrastive learning for video representation. *Computer Vision and Image Understanding*, page 103406, 2022. 1, 3, 4, 5, 6, 7
- [19] Ishan Dave, Zacchaeus Scheffer, Akash Kumar, Sarah Shiraz, Yogesh Singh Rawat, and Mubarak Shah. Gabriellav2: Towards better generalization in surveillance videos for action detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, pages 122–132, January 2022. 1
- [20] Ishan Rajendrakumar Dave, Chen Chen, and Mubarak Shah. Spact: Self-supervised privacy preservation for action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20164–20173, 2022. 3
- [21] Ali Diba, Mohsen Fayyaz, Vivek Sharma, Manohar Paluri, Jürgen Gall, Rainer Stiefelhagen, and Luc Van Gool. Large scale holistic video understanding. In *European Conference on Computer Vision*, pages 593–610. Springer, 2020. 1
- [22] Haoqi Fan, Bo Xiong, Karttkeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6824–6835, 2021. 3
- [23] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. In *Advances in Neural Information Processing Systems*, 2022. 3, 8
- [24] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 6202–6211, 2019. 6
- [25] Christoph Feichtenhofer, Haoqi Fan, Bo Xiong, Ross Girshick, and Kaiming He. A large-scale study on unsupervised spatiotemporal representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3299–3309, June 2021. 3

- [26] A. Gammerman, V. Vovk, and V. Vapnik. Learning by transduction. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, UAI'98, page 148–155, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc. [2](#)
- [27] Kirill Gavrilyuk, Mihir Jain, Iliia Karmanov, and Cees GM Snoek. Motion-augmented self-training for video recognition at smaller scale. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10429–10438, 2021. [6](#), [7](#)
- [28] Rohit Girdhar, Du Tran, Lorenzo Torresani, and Deva Ramanan. Distinit: Learning video representations without a single labeled video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 852–861, 2019. [6](#)
- [29] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33:21271–21284, 2020. [1](#)
- [30] Rohit Gupta, Naveed Akhtar, Ajmal Mian, and Mubarak Shah. On higher adversarial susceptibility of contrastive self-supervised learning. *arXiv preprint arXiv:2207.10862*, 2022. [3](#)
- [31] Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 312–329. Springer, 2020. [6](#), [7](#)
- [32] K. Hara, H. Kataoka, and Y. Satoh. Towards good practice for action recognition with spatiotemporal 3d convolutions. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 2516–2521, 2018. [3](#)
- [33] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. [3](#)
- [34] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. [3](#)
- [35] Neziha Jaouedi, Noureddine Boujnah, Oumayma Htiwich, and Med Salim Bouhlel. Human action recognition to human behavior analysis. In *2016 7th International Conference on Sciences of Electronics, Technologies of Information and Telecommunications (SETIT)*, pages 263–266. IEEE, 2016. [1](#)
- [36] Simon Jenni and Hailin Jin. Time-equivariant contrastive video representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9970–9980, 2021. [1](#), [3](#)
- [37] Simon Jenni, Givi Meishvili, and Paolo Favaro. Video representation learning by recognizing temporal transformations. In *The European Conference on Computer Vision (ECCV)*, August 2020. [3](#)
- [38] Longlong Jing, Toufiq Parag, Zhe Wu, Yingli Tian, and Hongcheng Wang. Videoss!: Semi-supervised learning for video classification. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1110–1119, 2021. [6](#), [7](#)
- [39] Longlong Jing, Xiaodong Yang, Jingen Liu, and Yingli Tian. Self-supervised spatiotemporal feature learning via video rotation prediction. *arXiv preprint arXiv:1811.11387*, 2018. [6](#), [7](#)
- [40] Thorsten Joachims. Transductive inference for text classification using support vector machines. In *Icml*, volume 99, pages 200–209, 1999. [2](#)
- [41] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*, pages 3581–3589, 2014. [2](#)
- [42] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: a large video database for human motion recognition. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2011. [6](#)
- [43] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR (Poster)*. OpenReview.net, 2017. [2](#)
- [44] Dong-Hyun Lee. Pseudo-label : The simple and efficient semi-supervised learning method for deep neural networks. 2013. [2](#)
- [45] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896, 2013. [6](#)
- [46] Yixuan Li, Lei Chen, Runyu He, Zhenzhi Wang, Gangshan Wu, and Limin Wang. Multisports: A multi-person video dataset of spatio-temporally localized sports actions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13536–13545, 2021. [1](#)
- [47] Yanghao Li, Chao-Yuan Wu, Haoqi Fan, Karttikeya Mangalam, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Improved multiscale vision transformers for classification and detection. *arXiv preprint arXiv:2112.01526*, 2021. [3](#)
- [48] Bin Liu, Zhirong Wu, Han Hu, and Stephen Lin. Deep metric transfer for label propagation with limited annotated data. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. [2](#)
- [49] Jixin Liu, Rong Tan, Guang Han, Ning Sun, and Sam Kwong. Privacy-preserving in-home fall detection using visual shielding sensing and private information-embedding. *IEEE Transactions on Multimedia*, 2020. [1](#)
- [50] Wenhe Liu, Guoliang Kang, Po-Yao Huang, Xiaojun Chang, Yijun Qian, Junwei Liang, Liangeke Gui, Jing Wen, and Peng Chen. Argus: Efficient activity detection system for extended video analysis. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, March 2020. [1](#)
- [51] Fuchen Long, Zhaofan Qiu, Yingwei Pan, Ting Yao, Jiebo Luo, and Tao Mei. Stand-alone inter-frame attention in video models. In *Proceedings of the IEEE/CVF Conference*

- on *Computer Vision and Pattern Recognition*, pages 3192–3201, 2022. 3
- [52] 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*, 41:1979–1993, 2018. 2
- [53] Tian Pan, Yibing Song, Tianyu Yang, Wenhao Jiang, and Wei Liu. Videomoco: Contrastive video representation learning with temporally adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11205–11214, 2021. 1, 3
- [54] Yunchen Pu, Zhe Gan, Ricardo Henao, Xin Yuan, Chunyuan Li, Andrew Stevens, and Lawrence Carin. Variational autoencoder for deep learning of images, labels and captions. In *Advances in neural information processing systems*, pages 2352–2360, 2016. 2
- [55] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6964–6974, 2021. 1, 3, 5
- [56] Mamshad Nayeem Rizve, Ugur Demir, Praveen Tirupattur, Aayush Jung Rana, Kevin Duarte, Ishan R Dave, Yogesh S Rawat, and Mubarak Shah. Gabriella: An online system for real-time activity detection in untrimmed security videos. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 4237–4244. IEEE, 2021. 1
- [57] Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *International Conference on Learning Representations*, 2021. 1, 2, 6
- [58] Mamshad Nayeem Rizve, Navid Kardan, Salman Khan, Fahad Shahbaz Khan, and Mubarak Shah. Openldn: Learning to discover novel classes for open-world semi-supervised learning. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*, pages 382–401. Springer, 2022. 2
- [59] Mamshad Nayeem Rizve, Navid Kardan, and Mubarak Shah. Towards realistic semi-supervised learning. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*, pages 437–455. Springer, 2022. 2
- [60] Michael Ryoo, AJ Piergiovanni, Anurag Arnab, Mostafa Dehghani, and Anelia Angelova. Tokenlearner: Adaptive space-time tokenization for videos. *Advances in Neural Information Processing Systems*, 34:12786–12797, 2021. 3
- [61] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 1163–1171. Curran Associates, Inc., 2016. 2
- [62] Madeline C. Schiappa, Yogesh S. Rawat, and Mubarak Shah. Self-supervised learning for videos: A survey. *ACM Computing Surveys*, dec 2022. 3
- [63] Weiwei Shi, Yihong Gong, Chris Ding, Zhiheng MaXiaoyu Tao, and Nanning Zheng. Transductive semi-supervised deep learning using min-max features. In *The European Conference on Computer Vision (ECCV)*, September 2018. 1, 2
- [64] Ankit Singh, Omprakash Chakraborty, Ashutosh Varshney, Rameswar Panda, Rogerio Feris, Kate Saenko, and Abir Das. Semi-supervised action recognition with temporal contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10389–10399, 2021. 1, 3, 6
- [65] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 2, 3, 6
- [66] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 6
- [67] Li Tao, Xueting Wang, and Toshihiko Yamasaki. Self-supervised video representation learning using inter-intra contrastive framework. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2193–2201, 2020. 3
- [68] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. 1, 2, 6
- [69] Hayato Terao, Wataru Noguchi, Hiroyuki Iizuka, and Masahito Yamamoto. Compressed video ensemble based pseudo-labeling for semi-supervised action recognition. *Machine Learning with Applications*, page 100336, 2022. 3
- [70] Anyang Tong, Chao Tang, and Wenjian Wang. Semi-supervised action recognition from temporal augmentation using curriculum learning. *IEEE Transactions on Circuits and Systems for Video Technology*, 2022. 1, 3, 6, 7
- [71] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. VideoMAE: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In *Advances in Neural Information Processing Systems*, 2022. 3, 8
- [72] 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. 3
- [73] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018. 6
- [74] Jue Wang, Gedas Bertasius, Du Tran, and Lorenzo Torresani. Long-short temporal contrastive learning of video transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14010–14020, 2022. 3
- [75] Jinpeng Wang, Yuting Gao, Ke Li, Xinyang Jiang, Xiaowei Guo, Rongrong Ji, and Xing Sun. Enhancing unsupervised

- video representation learning by decoupling the scene and the motion. In *The AAAI Conference on Artificial Intelligence (AAAI)*, 2021. 1, 3
- [76] Jiangliu Wang, Jianbo Jiao, and Yun-Hui Liu. Self-supervised video representation learning by pace prediction. In *The European Conference on Computer Vision (ECCV)*, August 2020. 3
- [77] Junfei Xiao, Longlong Jing, Lin Zhang, Ju He, Qi She, Zongwei Zhou, Alan Yuille, and Yingwei Li. Learning from temporal gradient for semi-supervised action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3252–3262, 2022. 1, 3, 6, 7
- [78] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33:6256–6268, 2020. 2
- [79] Bo Xiong, Haoqi Fan, Kristen Grauman, and Christoph Feichtenhofer. Multiview pseudo-labeling for semi-supervised learning from video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7209–7219, 2021. 1, 3, 6, 7
- [80] Yinghao Xu, Fangyun Wei, Xiao Sun, Ceyuan Yang, Yujun Shen, Bo Dai, Bolei Zhou, and Stephen Lin. Cross-model pseudo-labeling for semi-supervised action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2959–2968, 2022. 3, 6, 7
- [81] Ting Yao, Yiheng Zhang, Zhaofan Qiu, Yingwei Pan, and Tao Mei. Seco: Exploring sequence supervision for unsupervised representation learning. *arXiv preprint arXiv:2008.00975*, 2020. 3
- [82] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021. 3
- [83] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1476–1485, 2019. 3, 6
- [84] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems*, 34:18408–18419, 2021. 2
- [85] Chenyang Zhang, Yingli Tian, and Elizabeth Capezuti. Privacy preserving automatic fall detection for elderly using rgb-d cameras. In *International Conference on Computers for Handicapped Persons*, pages 625–633. Springer, 2012. 1
- [86] Yanyi Zhang, Xinyu Li, Chunhui Liu, Bing Shuai, Yi Zhu, Biagio Brattoli, Hao Chen, Ivan Marsic, and Joseph Tighe. Vidtr: Video transformer without convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13577–13587, 2021. 3
- [87] Hang Zhao, Antonio Torralba, Lorenzo Torresani, and Zhicheng Yan. Hacs: Human action clips and segments dataset for recognition and temporal localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8668–8678, 2019. 1
- [88] Yuliang Zou, Jinwoo Choi, Qitong Wang, and Jia-Bin Huang. Learning representational invariances for data-efficient action recognition. *arXiv preprint arXiv:2103.16565*, 2021. 6
- [89] Yuliang Zou, Jinwoo Choi, Qitong Wang, and Jia-Bin Huang. Learning representational invariances for data-efficient action recognition. *arXiv preprint arXiv:2103.16565*, 2021. 7