

# SpeedNet: Learning the Speediness in Videos

Sagie Benaim<sup>1,2\*</sup> Ariel Ephrat<sup>1</sup> Oran Lang<sup>1</sup> Inbar Mosseri<sup>1</sup> William T. Freeman<sup>1</sup>  
Michael Rubinstein<sup>1</sup> Michal Irani<sup>1,3</sup> Tali Dekel<sup>1</sup>

<sup>1</sup>Google Research <sup>2</sup>Tel Aviv University <sup>3</sup>Weizmann Institute

## Abstract

We wish to automatically predict the “speediness” of moving objects in videos—whether they move faster, at, or slower than their “natural” speed. The core component in our approach is SpeedNet—a novel deep network trained to detect if a video is playing at normal rate, or if it is sped up. SpeedNet is trained on a large corpus of natural videos in a self-supervised manner, without requiring any manual annotations. We show how this single, binary classification network can be used to detect arbitrary rates of speediness of objects. We demonstrate prediction results by SpeedNet on a wide range of videos containing complex natural motions, and examine the visual cues it utilizes for making those predictions. Importantly, we show that through predicting the speed of videos, the model learns a powerful and meaningful space-time representation that goes beyond simple motion cues. We demonstrate how those learned features can boost the performance of self-supervised action recognition, and can be used for video retrieval. Furthermore, we also apply SpeedNet for generating time-varying, adaptive video speedups, which can allow viewers to watch videos faster, but with less of the jittery, unnatural motions typical to videos that are sped up uniformly.

## 1. Introduction

A human observer can often easily notice if an object’s motion is sped up or slowed down. For example, if we play a video of a dancer at twice the speed ( $2\times$ ), we can notice unnatural, fast and jittery motions. In many cases, we have prior knowledge about the way objects move in the world (people, animals, cars); we know their typical dynamics and their natural *rate of motion*.

In this paper, we seek to study how well we can train a machine to learn such concepts and priors about objects’ motions. Solving this task requires high-level reasoning and understanding of the way different objects move in the world. We achieve this by training a single model, *SpeedNet*, to perform a basic binary classification task: estimate whether an object in an input video sequence is moving at its normal speed, or faster than the normal speed (Fig. 1, top). That is, given a set of  $L$  frames in an  $L$ -fps video as input, we set to predict whether those frames depict 1 second of the object’s motion in the world (normal speed), or

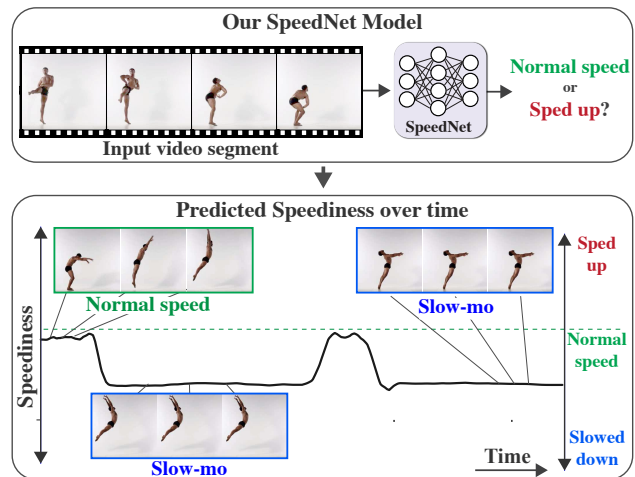


Figure 1. Given an input video, our method automatically predicts the “speediness” of objects in the video—whether they move faster, at, or slower than their natural speed. Bottom: a video of a dancer alternates between normal speed and slow motion playback, as correctly captured by our speediness prediction over time. Note that **speediness**  $\neq$  **magnitude of motions**, see Fig. 2. The core component of our method is SpeedNet (top)—a novel deep network that detects whether an object is moving at its normal speed, or faster than its normal speed.

more than 1 second (the object/video is sped up). We preferred this approach over directly predicting (regressing to) the playback rate of the video, because our ultimate goal is to determine whether motion in a given video is natural or not, a task for which a regression objective may be unnecessarily difficult to learn.

We dub this the “speediness” of the object. We then show how this basic binary classification model can be applied at test time to predict arbitrary rates of speediness in videos, when objects are sped-up, or slowed-down, by different factors (Fig. 1, bottom). The model is trained on Kinetics [16], a large corpus of natural videos of human actions, in a self-supervised manner, without requiring manual labels.

Our motivation for this study is twofold. First, we ask if it is possible to train a reliable classifier on a large-scale collection of videos to predict if an object’s motion is sped up or played at normal speed—what would such a model learn about the visual world in order to solve this task? Second, we show that a well-trained model for predicting speediness can support a variety of useful applications.

Training SpeedNet, however, is far from trivial. In hu-

\*Performed this work while an intern at Google.

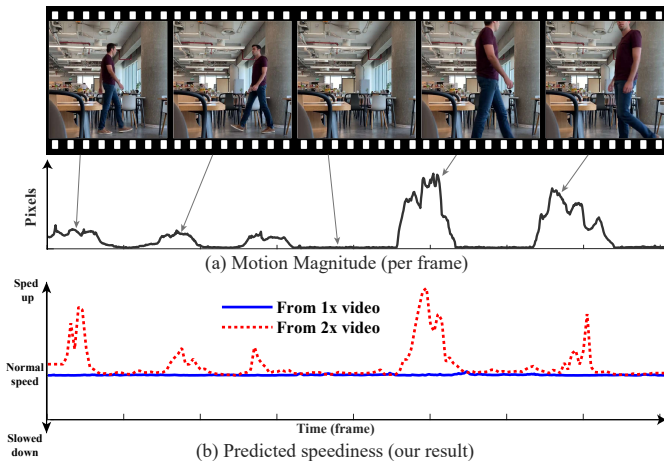


Figure 2. **Speediness  $\neq$  motion magnitude.** A person is walking back and forth, at first further away from the camera, then closer to the camera (top). The magnitude of motions varies significantly throughout the sequence (in particular, the motions get larger when the person is closer to the camera; middle plot), but our SpeedNet model is able to produce a stable classification of normal speed throughout the video (speediness score close to zero). If we input to SpeedNet the video played at twice the speed (2 $\times$ ), then the walking segments do indeed get recognized as faster-than-normal human motions (higher speediness scores), whereas the static segments (no person in the scene), are classified as normal speed.

mans, the ability to correctly classify an object’s speed continues to improve even throughout adolescence [23], implying that a developed mind is required to perform this task. In addition to the high-level reasoning necessary for solving this task, a major challenge in training a neural network to automatically predict speediness is to avoid their tendency to detect easy shortcuts, e.g., to rely on artificial, low-level cues, such as compression artifacts, whenever possible. This usually results in near-perfect accuracy on the learning task, which is an outcome we want to avoid: we seek a semantic, as opposed to artificial, understanding of speediness, by examining the *actual* motions, and not relying on artifacts (compression, aliasing) related to the way in which the sped-up videos are generated. We describe the strategies employed to mitigate artificial shortcuts in Sec. 3.1.

Another challenging aspect of training SpeedNet is to go beyond the trivial case of motion magnitude for determining speediness. Relying only on the speed of moving objects, using motion magnitude alone, would discriminate between, e.g. two people walking normally at two different distances from the camera (Fig. 2). We address the speediness prediction capability of optical flow in Sec. 5.1, and demonstrate the clear prediction superiority of SpeedNet over a naive flow-based baseline method. However, the correlation between speediness and motion magnitude does pose a formidable challenge to our method in cases of extreme camera or object motion.

We demonstrate our model’s speediness prediction results on a wide range of challenging videos, containing

complex natural motions, like dancing and sports. We visualize and examine the visual cues the model uses to make those predictions. We also apply the model for generating time-varying, adaptive speedup videos, speeding up objects more (or less) if their speediness score is low (or high), so that when sped up, their motion looks more natural to the viewer. This is in contrast to traditional video speedup, e.g. in online video streaming websites, which use uniform/constant speedup, producing unnatural, jittery motions. Given a desired global speedup factor, our method calculates a smooth curve for per-frame speedup factor based on the speediness score of each frame. The details of this algorithm are described in Sec. 4.2.

Lastly, we also show that through learning to classify the speed of a video, SpeedNet learns a powerful space-time representation that can be used for self-supervised action recognition and for video retrieval. Our action recognition results are competitive with the state-of-the-art self-supervised methods on two popular benchmarks, and beat all other methods which pre-train on Kinetics. We also show promising results on cross-video clip retrieval.

Our videos, results, and supplementary material are available on the project web page: <http://speednet-cvpr20.github.io>.

## 2. Related work

**Video playback speed classification.** Playback speed classification is considered a useful task in itself, especially in the context of sports broadcasts, where replays are often played at different speeds. A number of works try to detect replays in sports [36, 5, 14, 17]. However, these works usually employ a specific domain analysis, and use a supervised approach. Our method, however, works on any type of video and does not use any information unique to a specific sport. To the best of our knowledge, there exists no public dataset for detection of slow motion playback speed.

**Video time remapping.** Our variable speedup technique produces non-uniform temporal sampling of the video frames. This idea was explored by several papers. The early seminal work of Bennett and McMillan [2] calculates an optimal non-uniform sampling to satisfy various visual objectives, as captured by error metrics defined between pairs of frames. Zhou *et al.* [42] use a measure of “frame importance” based on motion saliency estimation, to select important frames. Petrovic *et al.* [29] perform query-based adaptive video speedup—frames similar to the query clip are played slower, and different ones faster. An important task is of intelligent fast-forwarding of regular videos [21], or egocentric videos [31, 33, 11, 32], where frames are selected to preserve the gist of the video, while allowing the user to view it in less time. All of these works try to select frames based on a saliency measure to keep the maximal “information” from the original video, or minimal camera jitter in the case of first-person videos. In contrast, our work focuses on detecting regions which are played at

slower than their natural speed. This allows for optimizing a varying playback rate, such that speedup artifacts of moving objects (as detected by the model) will be less noticeable.

**Self-supervised learning from videos.** Using video as a natural source of supervision has recently attracted much interest [15], and many different video properties have been used as supervision, such as: cycle consistency between video frames [37, 8]; distinguishing between a video frame sequence and a shuffled version of it [25, 10, 40]; solving space-time cubic puzzles [19]. Another common task is predicting the future, either by predicting pixels of future frames [24, 6], or an embedding of a future video segment [13]. Ng *et al.* [27] try to predict optical flow, and Vondrick *et al.* [35] use colorization as supervision.

The works most related to ours are those which try to predict the arrow of time [38, 30]. This task can be stated as classifying the playback speed of the video between  $-1$  and  $+1$ , as opposed to our work which attempts to discriminate between different positive video speeds. In concurrent work, Epstein *et al.* [9] leverage the intrinsic speed of video to predict unintentional action in videos.

### 3. SpeedNet

The core component of our method is *SpeedNet*—a deep neural network designed to determine whether objects in a video are moving at, or faster than, their normal speed. As proxies for natural and unnaturally fast movement, we train SpeedNet to discriminate between videos played at normal speed and videos played at twice ( $2\times$ ) their original speed. More formally, the learning task is: given a set of  $L$  frames extracted from an  $L$ -fps video as input, SpeedNet predicts if those frames contain 1 second of movement in the world (i.e., *normal speed*), or 2 seconds (i.e., *sped-up*).

It is important to note that videos played at twice the original speed do not always contain unnatural motion. For example, slow walking sped up to fast walking can still look natural. Similarly, when nothing is moving in the scene, a video played at  $2\times$  will still show no motion. Thus, the proxy task of discriminating between  $1\times$  and  $2\times$  speeds does not always accurately reflect our main objective of predicting speediness. Consequently, we do not expect (or desire for) our model to reach perfect accuracy. Moreover, this property of network ambiguity in cases of slow vs. fast natural motion is precisely what facilitates the downstream use of SpeedNet predictions to “gracefully” speed up videos.

We describe and demonstrate how we can use this model for predicting the speediness of objects in natural videos played at arbitrary speeds. The motivation for solving this binary classification problem rather than directly regressing to the video’s playback rate is because our goal is to determine whether or not the motion in a given video is natural, a task for which a regression objective may be unnecessarily difficult to learn. Moreover, discriminating between two different speeds is more natural for humans as well. We next describe the different components of our framework.

### 3.1. Data, supervision, and avoiding artificial cues

SpeedNet is trained in a self-supervised manner, without requiring any manually labeled videos. More specifically, our training and testing sets contain two versions of every video segment, a normal speed version, and a sped-up version constructed by temporally subsampling video frames.

Previous work has shown networks have a tendency to use shortcuts—artificial cues present in the training data, to help them solve the task at hand [38, 13, 7]. Our network too is susceptible to these cues, and we attempt to avoid potential shortcuts by employing the following strategies:

**Spatial augmentations.** Our base network, defined in Sec. 3.2, is fully convolutional, so its input can be of arbitrary dimensions. During training, we randomly resize the input video clip to a spatial dimension  $N$  of between 64 and 336 pixels. The blurring which occurs during the resize process can help mitigate potential pixel intensity jitter caused by MPEG or JPEG compression of each frame. After passing the input through the base network, we perform spatial global max pooling over the regions in the resulting space-time features. Since the input is of variable size, these regions correspond to differently sized regions in the original, unresized input. This forces our network not to rely only on size-dependent factors, such as motion magnitude.

**Temporal augmentations.** We would like to sample a video at either normal speed or twice its normal speed. To introduce variability in the time domain, for normal speed we sample frames at a rate of  $1\times-1.2\times$  and for the sped-up version we sample at  $1.7\times-2.2\times$ . In more detail, we choose  $3T$  consecutive frames from a given video. For normal speed, we randomly pick a skip factor  $f$  between  $1-1.2$  and skip frames with probability  $1-1/f$ . We then choose  $T$  consecutive frames from the remaining frames. For the sped-up version,  $f$  is chosen between  $1.7-2.2$ .

**Same-batch training.** For each clip (of  $3T$  consecutive frames), we construct a normal speed and a sped-up video, each of length  $T$ , in the manner described above. We train our model such that each batch contains both normal-speed and sped-up versions of each video clip. We found that this way, our network is significantly less reliant on artificial cues. We note the same type of training was found to be critical in other self-supervised works such as [12]. See Tab. 1 and the discussion in Sec. 5.1 for the quantitative effect of these augmentation strategies.

### 3.2. SpeedNet architecture

Our architecture is illustrated in Fig. 3. The input to the network is a  $T \times N \times N$  video segment, which is either sampled from a normal speed video, or its sped-up version ( $T$  and  $N$  denote the temporal and spatial dimensions, respectively). The input segment is then passed to a fully convolutional base network that learns space-time features. The dimensions of the output features are  $T \times \frac{N}{32} \times \frac{N}{32} \times 1024$ . That is, the spatial resolution is reduced by a factor of 32,

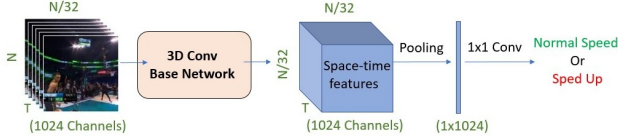


Figure 3. **SpeedNet architecture.** SpeedNet, the core model in our technique, is trained to classify an input video sequence as either normal speed, or sped up. Full details are provided in Sec. 3.2.

while the temporal dimension is preserved, and the number of channels is 1024. Our network architecture is largely based on S3D-G [39], a state-of-the-art action recognition model. There are two differences between our base model and the original S3D-G model: (i) In our model, temporal strides for all the max pooling layers are set to 1, to leave the input’s temporal dimension constant; (ii) We perform max spatial pooling and average temporal pooling over the resulting space-time features, as opposed to only average pooling in S3D-G.

We then collapse the temporal and spatial dimensions into a single channel. Intuitively, we want the prediction to be determined by the most dominant spatially moving object, whereas temporally, we would like to take into account the motion through the entire video segment to avoid sensitivity to instantaneous “spiky” motions. We therefore reduce the spatial dimension by applying global max pooling, and reduce the temporal dimension by applying global average pooling. This results in a  $1024D$  vector, which is then mapped to the final logits by a  $1 \times 1$  convolution  $W$ . Our model is trained using a binary cross entropy loss.

## 4. Adaptive video speedup

We use our model to adaptively speed up a test video  $v$ . The idea is: we speed up a video. As long as the network thinks that a segment in the resulting video is *not* sped-up, we can keep on speeding up that video segment even further.

### 4.1. From predictions to speedup scores

Given an input video  $v$ , we first generate a set of sped-up videos,  $v_0, \dots, v_k$ , by sub-sampling  $v$  by an exponential factor of  $\{X^i\}_{i=0}^k$ , where  $v_0$  is the original video. We used  $X = 1.25$  and  $k = 10$  in our experiments. We feed each video  $v_i$  into SpeedNet in a sliding window fashion; the network’s prediction for each window is assigned to the middle frame. This results in a temporally-varying prediction curve,  $P_{v_i}(t)$ , for each video  $v_i$ . Here,  $P_{v_i}(t)$  represents the (softmax) probability of normal speed. That is,  $P_{v_i}(t) \approx 1$  if SpeedNet’s prediction for a window centered at  $t$  is normal speed, and  $P_{v_i}(t) \approx 0$  if sped-up.

The set of speediness predictions  $\{P_{v_i}\}$  are first linearly interpolated (in time) to the temporal length of the longest curve ( $P_{v_0}$ ). We then binarize the predictions using a threshold  $\rho$  to obtain a *sped-up* or *not sped-up* classification per timestep, denoted by the set  $\{P_{v_i}^\rho\}$ . Each vector in this set of binary speediness predictions is then multiplied by its corresponding speedup factor  $X^i$  to obtain a set of

speedup vectors  $\{V_i(t)\}$ , which are combined into a single speedup vector  $V(t)$  by taking the maximum value at each timestep. In other words,  $V(t)$  contains the maximum possible speedup for each timestep that was still classified as *not sped-up*. The locally-adaptive speedups determine the *overall* speedup of the video that still seems “natural”.

## 4.2. Optimizing for adaptive speedup

Our main idea is to non-uniformly change the playback speed of a video based on its content. The motivation is similar to variable bitrate (VBR) encoding, where the bandwidth allocated to a data segment is determined by its complexity. The intuition is similar—some video segments, such as those with smooth, slow motion, can be sped up more than others without corrupting its “naturalness”.

How do we choose the threshold  $\rho$ , and how do we guarantee that we achieve the final *desired* overall speedup with least distortions? We test for nine thresholds:  $\rho \in \{0.1, \dots, 0.9\}$ , and select the one for which the overall speedup is closest to the desired one.

Given the per-frame speedup vector  $V(t)$ , as described above, our goal now is to estimate a smoothly varying speedup curve  $S(t)$ , which meets the user-given target speedup rate over the entire video. The motivation behind this process is that the speedup score of segments with little action will be high, meaning a human is less likely to notice the difference in the playback speed in those segments. We formulate this using the following objective:

$$\arg \min_S E_{\text{speed}}(S, V) + \beta E_{\text{rate}}(S, R_o) + \alpha E_{\text{smooth}}(S'),$$

where  $E_{\text{speed}}$  encourages speeding frames according to our estimated speedup score  $V$ .  $E_{\text{rate}}$  constrains the overall speedup over the entire video to match the user desired speedup rate  $R_o$ .  $E_{\text{smooth}}$  is a smoothness regularizer, where  $S'$  denotes the first derivative of  $S$ . We then plugin the optimal speedup  $S^*$  to adaptively play the video.

The graphs in Fig. 5 depict an example “*speediness curve*” (red), along with its corresponding final optimal speedup curve  $S^*$  (blue) for an *overall* target speedup of  $2 \times$ . We define a video’s “*speediness curve*” to be  $1 - \hat{V}(t)$ , where  $\hat{V}(t)$  is computed by normalizing  $V(t)$  to be in the range  $[0, 1]$ .

## 5. Experiments

For our experiments we use the Kinetics [16] dataset which consists of 246K train videos and 50K test clips played at 25 fps. As our method is self-supervised, we do not use the action recognition labels at training. We also test our model on the Need for Speed dataset (NFS) [18] which consists of 100 videos captured at 240 fps (a total of 380K frames). The dataset contains many different object actions such as moving balls, jumping people, skiing, and others. Testing our model’s performance on this dataset is important for evaluating our model’s generalization capability.

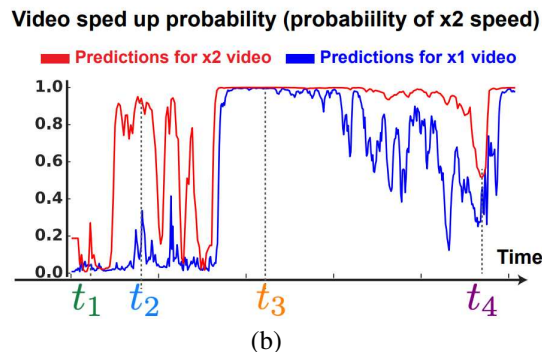
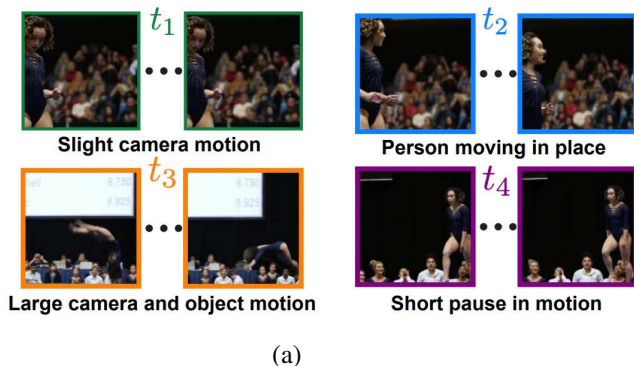


Figure 4. **Illustration of the network’s predictions.** (a) different segments from the same gymnast video are shown. (b) softmax probabilities of being “sped-up” (y axis) are displayed for the normal speed gymnast video (blue curve) and for the sped-up gymnast video (red curve). The segments shown in (a) are positioned in the graph. See further details in Sec. 5.1.1.

Batch	Model Type		Accuracy	
	Temporal	Spatial	Kinetics	NFS
Yes	Yes	Yes	75.6%	73.6%
No	Yes	Yes	88.2%	59.3%
No	No	Yes	90.0%	57.7%
No	No	No	96.9%	57.4%
Mean Flow			55.8%	55.0%

Table 1. **Ablation study.** We consider the effect of spatial and temporal augmentations described in Sec. 3 on the accuracy of SpeedNet for both Kinetics and NFS datasets. We also consider the effect of same-batch training (“Batch” in the table) vs. training only with random normal-speed and sped-up video clips in the same batch (see Sec. 3.1). In the last row, we consider the efficacy of training a simple network with only the mean flow magnitude of each frame.

## 5.1. SpeedNet performance

We assess the performance of SpeedNet (its ability to tell if a video is sped-up or not) on the test set of Kinetics and videos from the NFS dataset. We consider segments of 16 frames either played at normal speed (unchanged) or uniformly sped-up. For the NFS dataset, videos are uniformly sped-up up by  $10\times$  or  $20\times$  to give an effective frame rate of 24 fps (normal speed) or 12 fps (sped-up). As the frame rate of Kinetics videos is 25 fps, we expect these speedups to correspond to  $1\times$  and  $2\times$  speedups of Kinetics videos. The slight variation in frame rate is important for assessing the performance of our model on videos played at slightly different frame rates than trained on. At test time, we resize frames to a height of 224 keeping the aspect ratio of the original video, and then apply a  $224 \times 224$  center crop. No temporal or spatial augmentation is applied.

In Tab. 1, we consider the effect of training with or without: (1) temporal augmentations (2) spatial augmentations and (3) same-batch training (see Sec. 3.1). When not training in “same-batch training” mode, each batch consists of random normal-speed and sped-up video clips. When training SpeedNet without (1), (2) and (3), SpeedNet relies on learned “shortcuts”—artificial cues that are present in Kinetics, which lead to a misleading high test accuracy. However,

when tested on NFS, such cues are not present, and so the test accuracy drops to 57.4%. When using (1), (2) and (3), reliance on artificial cues is reduced significantly (accuracy drops from 96.9% to 75.6%), and the gap between the test accuracy of Kinetics and that of NFS drops to 2%, indicating better generalization. While chance level is 50%, recall that we do not expect SpeedNet to achieve accuracy close to 100%, as in many cases one cannot really tell if a video is sped-up or not (e.g., when there is no motion in the clip).

### 5.1.1 Prediction curves

Fig. 4 illustrates the predictions for a normal speed and a sped-up version of a gymnast video (more prediction results are on our `project` page). The predictions on the video played at normal speed ( $1\times$ ) are shown in blue and for sped-up ( $2\times$ ) in red. For a frame  $t$ , the predictions shown are  $1 - P_{v_0}(t)$  ( $1\times$ ) and  $1 - P_{v_3}(t)$  ( $2\times$ ), as detailed in Sec. 4. In particular, the predictions for  $2\times$  are linearly interpolated so as to be displayed on the same temporal axis. As can be seen, the region of slight camera motion ( $t_1$ ) is determined to be of normal speed for both  $1\times$  and  $2\times$ . A person moving in place ( $t_2$ ) is determined to be sped-up for  $2\times$  and of normal speed for  $1\times$ . Large camera and object motion ( $t_3$ ) is determined to be sped-up for both  $1\times$  and  $2\times$ . Lastly, a short pause in motion ( $t_4$ ), has roughly equal probability of being sped-up and of being normal speed for both  $1\times$  and  $2\times$ .

### 5.1.2 Comparison to optical flow

We consider the efficacy of training a baseline model whose input is the per-frame average flow magnitude for each example in our Kinetics training set. This results in a vector of size  $T$  for each video clip. We train a simple network with two fully connected layers, *ReLU* activations and batch normalization. As can be seen in Tab. 1, this model achieves only 55% accuracy on the test sets of Kinetics and NFS. A major limitation of the mean optical flow is its correlation with the distance of the object from the camera, which can be seen in Fig. 2. While SpeedNet is clearly superior compared to the flow baseline, it does tend to fail in scenarios which contain either extreme camera motion or very large object motion, such as fast moving objects and

motion very close to the camera. We hypothesize that this is because our training set does not contain enough normal-speed videos with such large frame-to-frame displacements, which is usually characteristic of videos being played at twice their original speed.

## 5.2. Generalization to arbitrary speediness rates

We tested our model on a variety of real-world videos downloaded from the web which contain slow motion effects, involving natural camera motion as well as complex human actions, including ballet dancing, Olympic gymnastics, skiing, and many more. Our algorithm manages to accurately predict which segments are at normal speed and which are *slowed down* by using the method described in Sec. 4. To emphasize, even though our SpeedNet model was trained on a dataset of videos at normal speed and 2× speed, we can use it within our framework to classify video segments which contain slow-motion, i.e. whose duration at playback is slower than real-time. A video clip is determined as “slow-motion” if its sped-up version is detected as “normal-speed”, such as the example shown in Fig. 1.

## 5.3. Adaptive speedup of real-world videos

To evaluate our adaptive speedup against that of uniform speedup, we seek videos where there is a large difference in the “speediness” of objects within the video. For example, for a 100m sprint run, the runner is initially walking towards the running blocks, then not moving at all just before the sprint (when at the blocks), and finally sprinting. We performed adaptive speedup on five such videos from YouTube, and then conducted a user study to determine the objective quality of our results. For each video, our adaptive speedup and its corresponding uniformly sped-up version are shown to the user at random, who is asked to select the sped-up version which “looks better”.

We conducted the study on 30 users with different research backgrounds, and for all five videos we presented, our adaptive speedup was preferred by a clear margin over uniform speedup, as shown in Fig. 6. An example of one of the adaptively sped-up videos used in our study is shown in Fig. 5, and all five videos are on our `project` page.

## 5.4. SpeedNet for self-supervised tasks

Solving the speediness task requires high-level reasoning about the natural movements of objects, as well as understanding of lower-level motion cues. Since SpeedNet is self-supervised, we evaluate the effectiveness of its internal representation on the self-supervised tasks of pre-training for action recognition and video retrieval.

### 5.4.1 Action recognition

Utilizing self-supervised pre-training to initialize action recognition models is an established and effective way for evaluating the internal representation learned via the self-supervised task. A good initialization is important especially when training on small action recognition datasets,

Method	Initialization	Supervised accuracy	
	Architecture	UCF101	HMDB51
Random init	S3D-G	73.8	46.4
ImageNet inflated	S3D-G	86.6	57.7
Kinetics supervised	S3D-G	96.8	74.5
CubicPuzzle [19]	3D-ResNet18	65.8	33.7
Order [40]	R(2+1)D	72.4	30.9
DPC [13]	3D-ResNet34	75.7	35.7
AoT [38]	T-CAM	79.4	-
SpeedNet (Ours)	S3D-G	<b>81.1</b>	<b>48.8</b>
Random init	I3D	47.9	29.6
SpeedNet (Ours)	I3D	66.7	43.7

Table 2. **Self-supervised action recognition.** Comparison of self-supervised methods on UCF101 and HMDB51 split-1. The top methods are baseline S3D-G models trained using various forms of initialization. All of the methods in the middle were trained with a self-supervised method on Kinetics and then fine-tuned on UCF101 and HMDB51. On the bottom, we show for reference our random init and SpeedNet accuracy when trained on I3D network.

such as UCF101 [34] and HMDB51 [20], as the generalization capability of powerful networks can easily be inhibited by quickly overfitting to the training set.

Fine-tuning a pre-trained SpeedNet model on either UCF101 or HMDB51 significantly boosts the action recognition accuracy over random initialization, which indicates that solving the speediness task led to a useful general internal representation. In Tab. 2 we show that our action recognition accuracy beats all other models pre-trained in a self-supervised manner on Kinetics. For reference, we include the performance of S3D-G network when pre-trained with ImageNet labels (ImageNet inflated), and when using the full supervision of Kinetics (Kinetics supervised). Both networks use additional supervision, which we do not.

The best performing self-supervised model we are aware of, DynamoNet [6], was pre-trained on YouTube-8M dataset [1], which is 10× larger than Kinetics, and whose raw videos are not readily available publicly. DynamoNet attains accuracies of 88.1% and 59.9% on UCF101 and HMDB51, respectively.

Note that our strong random init baselines for S3D-G are in part due to using 64 frames during training. SpeedNet was designed and trained for the specific requirements of “speediness” prediction and, as such, was not optimized for action recognition. For reference, when trained with a weaker architecture such as I3D [4], our speediness prediction drops to 63.1%, but we observe a larger absolute and relative improvement over the random init baseline for both datasets, as reported in Tab. 2.

### 5.4.2 Nearest neighbor retrieval

Another way of assessing the power of SpeedNet’s learned representation is by extracting video clip embeddings from the model, and searching for nearest neighbors in embeddings space. In particular, given a video clip of arbitrary

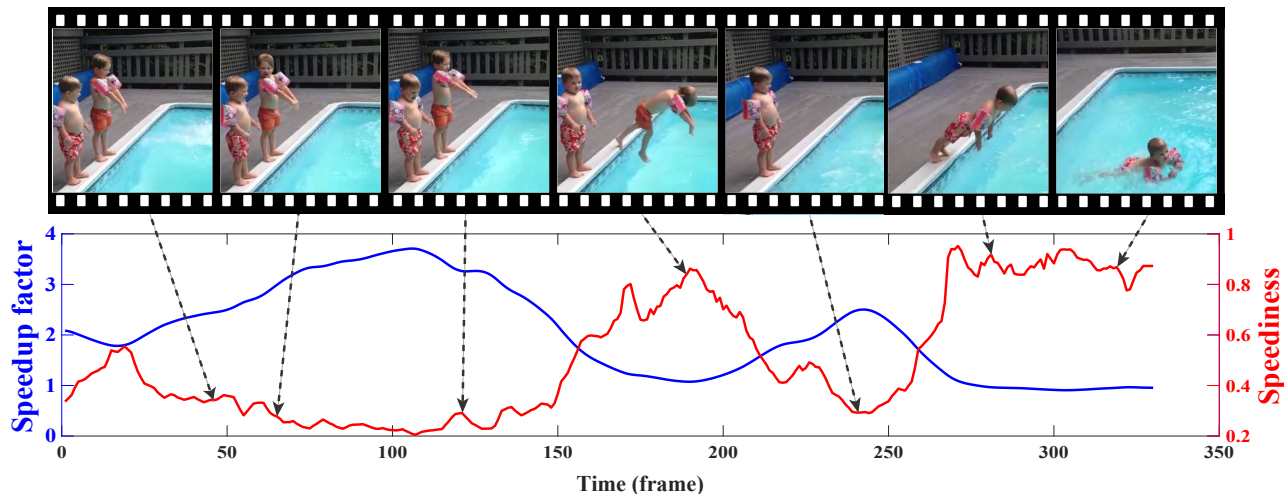


Figure 5. **Adaptive video speedup.** We apply the SpeedNet model for generating time-varying, adaptive speedup videos, based on the frames’ speediness curve (Sec. 4). Here we show the speediness curve and our resulting adaptive speedup factor for a video of two kids jumping into a pool. Several selected frames are shown at the top, pointing to their corresponding times within the sequence on the predicted speediness curve.

100m	61.5%	38.5%
Pool	77.8%	22.2%
High Jump	70.4%	29.6%
Dancing	81.5%	18.5%
Floor is Lava	59.3%	37%

■ Adaptive speedup (ours) 
 ■ Constant speedup 
 ■ Can’t tell. They look the same

Figure 6. **Adaptive video speedup user study.** We asked 30 participants to compare our adaptive speedup results with constant uniform speedup for 5 videos (without saying which is which), and select the one they liked better. Our adaptive speedup results were consistently (and clearly) preferred over uniform speedup.

spatial dimension and temporal length, we propose to use the max and average-pooled space-time activations, described in Sec. 3.2, as a  $1024D$  feature vector representing the clip. The experiments described in this section demonstrate that the extracted features encapsulate motion patterns in a way that facilitates the retrieval of other clips with similar behavior.

In this experiment, we extract a 16-frame sequence from a video (*query*, see Fig. 7), and aim to retrieve similar clips from either the same (longer) video (“within a video”), or from a collection of short video clips (“across videos”). For the former, we first extract the query feature vector from SpeedNet, and proceed to calculate feature vectors from the target video in sliding-window fashion over 16-frame windows. We then calculate a similarity graph by computing the cosine similarity score between the query feature and each of the target video features. In the first experiment, the query is of a basketball player shooting a 3-point shot, and similar clips of a different player are retrieved from further along in the same video, filmed from a slightly different angle and scale. In Fig. 7 (a) we show a representative frame from each peak in the similarity graph.

In the second experiment (Fig. 7 (b)), the query clip is

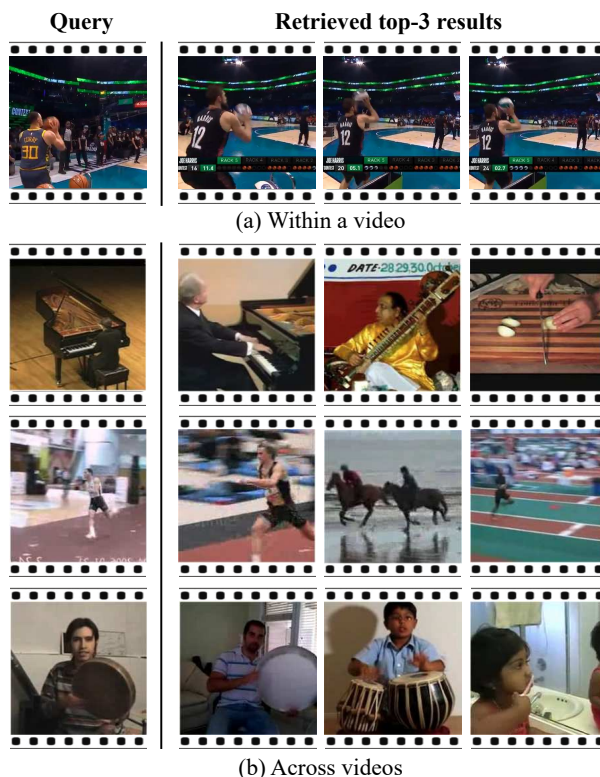


Figure 7. **Video clip retrieval.** The left column shows an image from a query clip, and the right three columns show the clips with the closest embeddings. In (a), the retrieval is from clips taken from further along in the same video. In (b), the results are retrieved from the entire UCF101 train set. Note that the embeddings focus more on the type of movement than the action class, for example the girl in the last row is making similar back/forth/up/down motions with her hand as the hand drummer in the query.

taken from the test set of UCF101, and we search for the nearest neighbors in the training set, again using cosine sim-

Method	Architecture	1	5	10	20	50
Jigsaw [28]	CFN	19.7	28.5	33.5	40.0	49.4
OPN [22]	OPN	19.9	28.7	34.0	40.6	51.6
Buchler [3]	CaffeNet	25.7	36.2	42.2	49.2	59.5
Order [40]	C3D	12.5	29.0	39.0	50.6	66.9
Order [40]	R(2+1)D	10.7	25.9	35.4	47.3	63.9
Order [40]	R3D	14.1	30.3	40.0	51.1	66.5
Ours	S3D-G	13.0	28.1	37.5	49.5	65.0

Table 3. **Recall-at-topK.** Top-K accuracy for different values of K for UCF101.

ilarity. SpeedNet mostly focuses on the type and speed of an object’s behavior, which is not always equivalent to the class of the video. For example, in Fig. 7, the girl in the last row is making similar back/forth/up/down motions with her hand as the hand drummer in the query.

We would, however, like to measure how correlated our learned representations are to specific action classes. We consider a third experiment where we measure the Recall-at-topK: a test clip is considered correctly classified if it is equivalent to the class of one of the  $K$  nearest training clips. We use the protocol of Xu *et al.* [40] (denoted *Order*). As can be seen in Tab. 3, our method is competitive with other self-supervised methods and only slightly inferior to [40].

### 5.5. Visualizing salient space-time regions

To gain a better understanding about which space-time regions contribute to our predictions, we follow the Class-Activation Map (CAM) technique [41] to visualize the energy of the last 3D layer, before the global max and average pooling (see Fig. 3). More specifically, we extract a  $T \times N \times N \times 1024$  feature map, where  $T$  and  $N$  are the temporal and spatial dimensions, respectively. We first reduce the number of channels to  $T \times N \times N$  using  $W$  (the weights that map from the  $1024D$  vector to the final logits). We then take the absolute value of the activation maps and normalize them between 0 and 1.

Fig. 8 depicts computed heat maps superimposed over sample frames from several videos are shown in . These examples portray a strong correlation between highly activated regions and the dominant mover in the scene, even when performing complex actions such as flips and articulated motions. For example, in the top row, second frame, the network attends to the motion of the leg, and in the second row, activations are highly concentrated on the body motion of the gymnast. Interestingly, the model is able to pick up the salient mover in the presence of significant camera motion.

In Fig. 9, we consider the video “Memory Eleven” [26], in which part of the frame is played in slow motion while the other part is played at normal speed. We use a similar visualization as in Fig. 8, but do not take the absolute value of the activation maps. While in Fig. 8 we are interested in overall important spatio-temporal regions for classification, in Fig. 9 we are interested in distinguishing between areas

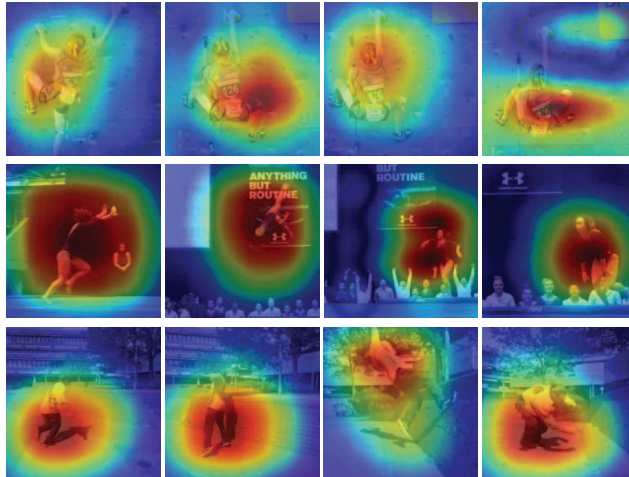


Figure 8. **Which space-time regions contribute the most to our speediness predictions?** CAM visualizations as detailed in Sec. 5.5. We visualize such regions as overlaid heat-maps, where red and blue correspond to high and low activated regions, respectively. Interestingly, the model is able to pick up the salient mover in the presence of significant camera motion.



Figure 9. **Spatially-varying speediness.** In the video “Memory Eleven” [26], one part of the frame is played in slow motion while the other part is played at regular speed. Using CAM visualizations without taking the absolute value (thus maintaining the direction of the activations, see Sec. 5.5), we can see that the model accurately predicts the speediness of each part of the frame part, from blue (normal speed), to red (slow motion).

in the frame used for classifying the video as normal and as sped-up. SpeedNet accurately predicts the speediness of each part, from blue (normal speed), to red (slow motion).

## 6. Conclusion

Our work studies the extent to which a machine can learn the “speediness” of moving objects in videos: whether an object is moving slower, at, or faster than its natural speed. To this end, we proposed SpeedNet, a model trained in a self-supervised manner to determine if a given video is being played at normal or twice its original speed. We showed that our model learns high level object motion priors which are more sophisticated than motion magnitude, and demonstrated the effectiveness of our model for several tasks: adaptively speeding up a video more “naturally” than uniform speedup; as self-supervised pre-training for action recognition; as a feature extractor for video clip retrieval.



## References

- [1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Apostol (Paul) Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. In *arXiv:1609.08675*, 2016.
- [2] Eric P Bennett and Leonard McMillan. Computational time-lapse video. In *ACM Transactions on Graphics (TOG)*, volume 26, page 102. ACM, 2007.
- [3] Uta Büchler, Biagio Brattoli, and Björn Ommer. Improving spatiotemporal self-supervision by deep reinforcement learning. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 797–814, Cham, 2018. Springer International Publishing.
- [4] J. Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. pages 4724–4733, 07 2017.
- [5] Chun-Min Chen and Ling-Hwei Chen. A novel method for slow motion replay detection in broadcast basketball video. *Multimedia Tools and Applications*, 74(21):9573–9593, 2015.
- [6] Ali Diba, Vivek Sharma, Luc Van Gool, and Rainer Stiefelhagen. Dynamonet: Dynamic action and motion network. *arXiv preprint arXiv:1904.11407*, 2019.
- [7] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1422–1430, 2015.
- [8] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. Temporal cycle-consistency learning. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [9] Dave Epstein, Boyuan Chen, and Carl Vondrick. Oops! predicting unintentional action in video. *arXiv preprint arXiv:1911.11206*, 2019.
- [10] 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.
- [11] Vinicius S Furlan, Ruzena Bajcsy, and Erickson R Nascimento. Fast forwarding egocentric videos by listening and watching. *arXiv preprint arXiv:1806.04620*, 2018.
- [12] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018.
- [13] 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.
- [14] Ali Javed, Khalid Bashir Bajwa, Hafiz Malik, and Aun Irtaza. An efficient framework for automatic highlights generation from sports videos. *IEEE Signal Processing Letters*, 23(7):954–958, 2016.
- [15] Longlong Jing and Yingli Tian. Self-supervised visual feature learning with deep neural networks: A survey. *arXiv preprint arXiv:1902.06162*, 2019.
- [16] 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.
- [17] Vahid Kiani and Hamid Reza Pourreza. An effective slow-motion detection approach for compressed soccer videos. *ISRN Machine Vision*, 2012, 2012.
- [18] Hamed Kiani Galoogahi, Ashton Fagg, Chen Huang, Deva Ramanan, and Simon Lucey. Need for speed: A benchmark for higher frame rate object tracking. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1125–1134, 2017.
- [19] Dahun Kim, Donghyeon Cho, and In So Kweon. Self-supervised video representation learning with space-time cubic puzzles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8545–8552, 2019.
- [20] 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.
- [21] Shuyue Lan, Rameswar Panda, Qi Zhu, and Amit K Roy-Chowdhury. Ffnet: Video fast-forwarding via reinforcement learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6771–6780, 2018.
- [22] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. pages 667–676, 10 2017.
- [23] Catherine Manning, David Aagten-Murphy, and Elizabeth Pellicano. The development of speed discrimination abilities. *Vision Research*, 70:27–33, 2012.
- [24] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015.
- [25] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision*, pages 527–544. Springer, 2016.
- [26] Bill Newsinger. Memory eleven. <http://vimeo.com/29213923>. Accessed: 2019-11-15.
- [27] Joe Yue-Hei Ng, Jonghyun Choi, Jan Neumann, and Larry S Davis. Actionflownet: Learning motion representation for action recognition. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1616–1624. IEEE, 2018.
- [28] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 69–84, Cham, 2016. Springer International Publishing.
- [29] Nemanja Petrovic, Nebojsa Jojic, and Thomas Huang. Adaptive video fast forward. *Multimedia Tools and Applications*, 26:327–344, 01 2005.
- [30] Lyndsey C Pickup, Zheng Pan, Donglai Wei, YiChang Shih, Changshui Zhang, Andrew Zisserman, Bernhard Scholkopf, and William T Freeman. Seeing the arrow of time. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [31] Yair Poleg, Tavi Halperin, Chetan Arora, and Shmuel Peleg. Egosampling: Fast-forward and stereo for egocentric videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4768–4776, 2015.

- [32] Michel Silva, Washington Ramos, Joao Ferreira, Felipe Chamone, Mario Campos, and Erickson R Nascimento. A weighted sparse sampling and smoothing frame transition approach for semantic fast-forward first-person videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2383–2392, 2018.
- [33] Michel M Silva, Mario FM Campos, and Erickson R Nascimento. Semantic hyperlapse: a sparse coding-based and multi-importance approach for first-person videos. In *Anais Estendidos da XXXII Conference on Graphics, Patterns and Images*, pages 56–62. SBC, 2019.
- [34] 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.
- [35] Carl Vondrick, Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, and Kevin Murphy. Tracking emerges by coloring videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 391–408, 2018.
- [36] Lei Wang, Xu Liu, Steve Lin, Guangyou Xu, and Heung-Yeung Shum. Generic slow-motion replay detection in sports video. In *2004 International Conference on Image Processing, 2004. ICIP'04.*, volume 3, pages 1585–1588. IEEE, 2004.
- [37] Xiaolong Wang, Allan Jabri, and Alexei A Efros. Learning correspondence from the cycle-consistency of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2566–2576, 2019.
- [38] Donglai Wei, Joseph J Lim, Andrew Zisserman, and William T Freeman. Learning and using the arrow of time. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [39] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 305–321, 2018.
- [40] Dejing Xu, Jun Xiao, Zhou Zhao, Jian Shao, Di Xie, and Yueting Zhuang. Self-supervised spatiotemporal learning via video clip order prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10334–10343, 2019.
- [41] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.
- [42] Feng Zhou, Sing Bing Kang, and Michael F Cohen. Time-mapping using space-time saliency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3358–3365, 2014.