

Audio-Visual Contrastive Learning with Temporal Self-Supervision

Simon Jenni,¹ Alexander Black,² John Collomosse^{1,2}

¹ Adobe Research

² University of Surrey

jenni@adobe.com, alex.black@surrey.ac.uk, collomos@adobe.com

Abstract

We propose a self-supervised learning approach for videos that learns representations of both the RGB frames and the accompanying audio without human supervision. In contrast to images that capture the static scene appearance, videos also contain sound and temporal scene dynamics. To leverage the temporal and aural dimension inherent to videos, our method extends temporal self-supervision to the audio-visual setting and integrates it with multi-modal contrastive objectives. As temporal self-supervision, we pose playback speed and direction recognition in both modalities and propose intra- and inter-modal temporal ordering tasks. Furthermore, we design a novel contrastive objective in which the usual pairs are supplemented with additional sample-dependent positives and negatives sampled from the evolving feature space. In our model, we apply such losses among video clips and between videos and their temporally corresponding audio clips. We verify our model design in extensive ablation experiments and evaluate the video and audio representations in transfer experiments to action recognition and retrieval on UCF101 and HMDB51, audio classification on ESC50, and robust video fingerprinting on VGG-Sound, with state-of-the-art results.

Introduction

Videos provide a rich source of information for audio-visual learning. Besides static moments in time (single video frames), they also contain the scene dynamics (object motion) and often include the sounds of the environment and scene objects. It seems hopeless to learn general representations that capture this rich semantic information in videos, *i.e.*, their appearance, motions, and sounds from such high-dimensional data through sparse human supervision. Self-supervised learning (SSL) (Doersch, Gupta, and Efros 2015; Chen et al. 2020b; He et al. 2020) has emerged as a viable alternative to supervised learning in recent years. Such methods might be better suited for general video representation learning since they are not constrained by the prohibitive cost of exhaustive human annotations on video. However, since most current self-supervised methods are tailored to static images, they might not effectively use videos' added temporal and aural dimensions. A self-supervised learning task that successfully integrates the static scene appearance

and the aural and temporal features potentially results in a representation that better generalizes to downstream vision applications, such as action recognition, video retrieval, or robust video content fingerprinting.

Indeed, recent works that explored the aural and temporal dimensions of videos in isolation have demonstrated that they are both effective self-supervision signals. Several works (Morgado, Vasconcelos, and Misra 2021; Patrick et al. 2020; Alwassel et al. 2019) demonstrate that audio-visual contrastive learning often performs better than uni-modal contrastive learning (*i.e.*, using only the RGB frames). Likewise, temporal reasoning tasks (Misra, Zitnick, and Hebert 2016; Jenni and Jin 2021; Dave et al. 2021) have demonstrated good transfer performance, especially for downstream tasks where motion is the main discerning factor (as opposed to static scene appearance).

In contrast, our work aims to leverage both sound and time as learning signals in a unified model architecture and training objective. To this end, we extend temporal self-supervision to the audio domain and propose cross-modal audio-visual temporal reasoning tasks. Concretely, we pose playback-speed and -direction recognition (Wei et al. 2018; Benaim et al. 2020; Jenni, Meishvili, and Favaro 2020), as a pretext task for audio representation learning and propose temporal clip ordering as a task for both intra-modal (*e.g.*, audio-audio) and cross-modal (*e.g.*, audio-video) learning (see Figure 2). Furthermore, we introduce a model architecture and training objective for contrastive audio-visual learning that supplements these temporal learning tasks. Towards this goal, we carefully study how the inclusion and exclusion of different intra- and inter-modal contrastive objectives influences downstream performance. Our key findings for optimal audio-visual contrastive learning are 1. inclusion of video-video contrastive terms 2. temporally aligned cross-modal positives, and 3. exclusion of audio-audio contrastive terms (see Figure 1).

We further explore the design of the contrastive loss terms (Wu et al. 2018), *i.e.*, how to build positive and negative pairs for effective learning. In constructing our contrastive objective, we take inspiration from recent image-based methods (Dwibedi et al. 2021; Koohpayegani, Tejankar, and Pirsiavash 2021) and extend the set of positive samples with nearest neighbors in the evolving feature space. Thus, besides standard augmented views for positive sampling, we con-

sider nearest neighbors sampled from a queue of prior embeddings as additional positives. Notably, the neighborhood structure and sample weights are both calculated through cross-view similarity, *i.e.*, either through the feature space similarity to the augmented view (for intra-modal learning) or the temporally aligned sample from the other modality (for cross-modal learning). We also use this cross-view induced neighborhood structure to sample negative pairs in a sample-dependent manner. This allows us to control the difficulty of the negative samples, *e.g.*, preventing ambiguous or confusing negatives resulting from duplicates or heavy class imbalance, while also preventing possible collapse through the absence of any negatives.

We verify our model design in extensive ablation experiments and compare it to prior works in established action recognition and retrieval benchmarks on UCF101 and HMDB51. We also evaluate the audio branch of our model for environmental sound classification on ESC50. Finally, we demonstrate the effectiveness of fusing the learned audio-visual features for downstream video classification on Kinetics-600 and VGG-Sound, and for robust video content retrieval under novel content manipulations for video fingerprinting (Lee and Yoo 2008; Black et al. 2021) on VGG-Sound. We investigate video fingerprinting as a novel downstream application due to its growing importance given the ever-expanding scale of visual data online and the increasing threat and sophistication of malicious content manipulations.

Contributions. To summarize, we make the following contributions: 1) We introduce temporal self-supervision in the audio domain and the cross-modal setting; 2) We propose a contrastive loss design that extends the usual contrastive pairs with sample-dependent positives and negatives; 3) We explore various multi-modal contrastive model designs and demonstrate the importance of a) using temporally aligned positives for cross-modal terms and b) excluding audio-audio contrastive terms; 4) Finally, we demonstrate the quality of the learned audio-visual features in extensive transfer experiments to action recognition, video retrieval, audio classification, and a novel video fingerprinting benchmark.

Prior Work

Contrastive Video Representation Learning. Contrastive learning is arguably the most popular self-supervised learning approach in computer vision today. These methods are typically based on the task of discriminating training instances up to strong data-augmentation (Dosovitskiy et al. 2015; Wu et al. 2018), which was shown to be remarkably effective for unsupervised image representation learning (Chen et al. 2020b; He et al. 2020) and has inspired a line of novel self-supervised methods (Grill et al. 2020; Chen and He 2020; Caron et al. 2020; Wang, Liu, and Yu 2020). Recently, methods were proposed that extend the set of positive pairs with nearest neighbors in the learned embedding space (Dwivedi et al. 2021; Koohpayegani, Tejankar, and Pirsiavash 2021). Our loss design similarly uses the evolving feature space to extend the set of contrastive pairs. In contrast, our loss design retains the exact match, contains




Figure 1: **Illustration of Contrastive Loss Terms in our Model.** We demonstrate the main contrastive pairs in our formulation given an example video clip (yellow box in the middle) and its corresponding audio clip. Positives (solid green arrows) are constructed from differently augmented video clips of the same training instance (blue box) and *temporally aligned* pairs of the corresponding video and audio clips. Negatives (dashed red arrows) stem from m other video and audio clips from the current mini-batch or a memory bank of prior embeddings (gray box on the right). Additional positives from the memory bank are omitted from the figure. Note that our formulation does not contain any contrastive terms among audio clips.

multiple positives weighted based on cross-view similarity, and uses additional sample-dependent negatives.

Several recent works have explored contrastive learning on video. When dealing with video, the set of data augmentations can be extended with several temporal augmentations (*e.g.*, temporal crops). A natural extension is thus to add temporal augmentations to the set of data-augmentations that define the positive pairs for contrastive learning (Qian et al. 2020; Feichtenhofer et al. 2021). Other works instead propose to learn to discriminate among temporally augmented clips (Dave et al. 2021; Patrick et al. 2020), or learn to recognize the temporal input transformations in a multi-task approach (Bai et al. 2020; Jenni and Jin 2021). Our model combines contrastive learning among video clips with audio-visual contrastive and temporal self-supervised learning.

Temporal Self-Supervision. Classic self-supervised approaches were based on so-called pretext tasks. On images, popular examples are the ordering of image patches (Doersch, Gupta, and Efros 2015; Noroozi and Favaro 2016), the colorization of gray-scale images (Zhang, Isola, and Efros 2016, 2017), or the classification of sets of image transformations (Gidaris, Singh, and Komodakis 2018; Jenni and Favaro 2018; Jenni, Jin, and Favaro 2020). Pretext tasks that turned out particularly successful on video are based on recognizing temporal transformations. Some works explored the ordering of video frames (Misra, Zitnick, and Hebert 2016; Brattoli et al. 2017; Fernando et al. 2017; Lee et al. 2017) or whole video clips (Xu et al. 2019; Kim, Cho, and Kweon 2019), others the classification of the playback direction (Wei et al. 2018), the playback speed (Epstein, Chen, and Vondrick 2020; Benaim et al. 2020; Yao et al. 2020), or general temporal warpings (Jenni, Meishvili, and Favaro 2020). We also leverage temporal supervision and extend it to multi-modal audio-visual representation learning.




Figure 2: **Illustration of the Temporal Reasoning Tasks.** Besides contrastive terms, our model encompasses both per-clip classification tasks (blue arrows) about the playback-speed and -direction, and temporal ordering tasks (green arrows) which are performed both intra- and cross-modal (V: RGB frames, A: audio).

Audio-Visual Self-Supervised Learning. Another source of self-supervision on video can be found in the accompanying sound. Early works explored audio to learn single frame representations, *e.g.*, by predicting summary statistics of the sounds corresponding to a frame (Owens et al. 2016), or by recognizing if an audio snippet and image are temporally aligned (Owens and Efros 2018; Arandjelovic and Zisserman 2017). Similar to these image-based approaches (Korbar, Tran, and Torresani 2018) learned audio and video representations by recognizing when audio and video signals are synchronized. More recently, contrastive audio-visual learning for video achieved remarkable performance (Recasens et al. 2021). For example, (Alwassel et al. 2020) performs clustering in one domain (*e.g.*, audio) and uses the resulting clusters as supervision for the other domain (*e.g.*, video). (Morgado, Vasconcelos, and Misra 2021) demonstrate the effectiveness of cross-modal audio-visual contrastive learning and extend the set of positive samples within a modality with samples that show high cross-modal agreement. Other works even include language in audio-visual contrastive models (Alayrac et al. 2020; Akbari et al. 2021). We instead focus on audio-visual learning and propose incorporating temporal supervision in both modalities.

Model

Let $\mathcal{D}_v = \{v_1, v_2, \dots, v_N\}$ be a set of unlabelled training videos and let $\mathcal{D}_a = \{a_1, a_2, \dots, a_N\}$ be their corresponding audio tracks. Our goal is to learn a video encoder F_v (a 3D-ConvNet) and an audio encoder F_a (a 2D-ConvNet) without human supervision. The inputs to the two networks are assumed to be of shape $v_i \in \mathbb{R}^{T \times H \times W \times C}$ and $a_i \in \mathbb{R}^{f \times t}$, where a_i is a spectrogram representation of the audio track.

Temporal Input Augmentations. An essential component of modern SSL approaches is the set of data augmentations applied to the input. In contrastive learning, these input transformations define the set of learned invariances. They typically comprise color jittering and geometric transforma-

tions, like random resizing, cropping, and horizontal flipping. For our method, temporal transformations, *i.e.*, random temporal cropping and manipulations of the playback speed and direction, are particularly important. We will thus indicate the precise temporal manipulations with τ_r . Furthermore, we assume that τ_r has consistent behavior across modalities, *i.e.*, $\tau_r(v_j)$ and $\tau_r(a_r)$ represent the exact same moments in time for the audio and video domain.

Intra- and Inter-Modal Contrastive Learning

Our training objective comprises several predictive-contrastive loss terms. In general, we formulate these losses based on the two modalities involved and on the direction of the prediction, *e.g.*, indicating that the visual representation is being predicted from the audio. For the purpose of this discussion let $\nu_i^r = \psi_v(F_v(\tau_r(v_i)))$ denote the output of the video encoder followed by a projection MLP ψ_v for the input $\tau_r(v_i)$. Let similarly $\alpha_i^r = \psi_a(F_a(\tau_r(a_i)))$ be the feature vector for the corresponding audio track. Let further $\hat{\nu}_i^s$ be the feature of a different augmentation of the video v_i .

We illustrate the general form of the contrastive objective using the video-to-audio loss term, which is given by

$$\ell_{va}(\nu_i^r, \mathcal{P}_i^{va}, \mathcal{N}_i^{va}) = \sum_{p, w \in \mathcal{P}_i^{va}} -w \log \left(\frac{d(\phi_v(\nu_i^r), p)}{d(\phi_v(\nu_i^r), p) + \sum_{n \in \mathcal{N}_i^{va}} d(\phi_v(\nu_i^r), n)} \right), \quad (1)$$

where ϕ_v denotes a predictor MLP (following prior work (Grill et al. 2020)) and

$$d(x, y) := \exp \left(\frac{1}{\lambda} \frac{x^\top y}{\|x\|_2 \|y\|_2} \right), \quad (2)$$

is a measure of the similarity between the feature representations of x and y , and $\lambda = 0.2$ is a temperature parameter. Note that we do not back-propagate through the second argument y in Eq. 2. In this general formulation, the set \mathcal{P}_i defines the instance-dependent positive samples along with their weighting factor w , and \mathcal{N}_i defines the negatives for contrastive learning.

Sources for Positive and Negative Contrastive Samples.

We consider two sources for sampling the positive and negative pairs of the contrastive loss terms: 1. the set of examples in the mini-batch \mathcal{B} at each iteration, and 2. a memory bank of prior feature embeddings. In our model, we maintain a memory bank Q_v (implemented as a FIFO queue) for the video domain and a corresponding Q_a for audio. Let $|Q_v| = |Q_a| = n_q$ be the size of the memory banks and let $\text{NN}_{j:k}(\nu, Q_v)$ denote the sequence $\{\eta_j, \dots, \eta_k\}$ from the j -th to the k -th nearest-neighbor of ν in Q_v . For positive examples from the memory bank we further introduce a set of loss term weights $W_{1:k}(\nu) := \{\omega_1, \dots, \omega_k\}$, where each weight is given by

$$\omega_j := \frac{d(\nu, \eta_j)}{\sum_{\eta_l \in \text{NN}_{1:k}(\nu, Q_v)} d(\nu, \eta_l)}, \quad (3)$$

thus weighting each nearest neighbor proportional to their similarity to ν . The memory banks are updated with the

mean of the features from the two augmented views in each mini-batch, *i.e.*, $(\nu + \hat{\nu})/2$ in the case of Q_v .

We will now describe different instantiations of the contrastive losses and their positive and negative sample sets for the intra-modal and cross-modal objectives.

Visual-Visual Contrastive Term ℓ_{vv} . For a video feature vector ν_i^r in the case of video-video contrastive learning, we set $\mathcal{P}_i^{vv} = \{(\hat{\nu}_i^s, 1)\} \cup \text{NN}_{1:k}(\hat{\nu}_i^s, Q_v) \times W_{1:k}(\hat{\nu}_i^s)$, where $\text{NN}_{1:k}(\hat{\nu}_i^s, Q_v)$ is the set of the first k nearest neighbors of $\hat{\nu}_i^s$ extracted from Q_v . We set $k = 5$ in our experiments. The set of negatives is constructed as $\mathcal{N}_i^{vv} = \{\nu_j \in \mathcal{B} | j \neq i\} \cup \text{NN}_{q:q+m}(\hat{\nu}_i^s, Q_v)$ and contains all the video features not belonging to v_i that are in the current training mini-batch \mathcal{B} , as well as m additional negatives sampled from the memory queue as the q -th up to the $(q+m)$ -th nearest neighbor of $\hat{\nu}_i^s$. By default we set $q = \frac{n_a}{2}$, thus starting from the neighbor in Q_v with median distance to $\hat{\nu}_i^s$ and set $m = 2048$.

Audio-Visual Contrastive Terms ℓ_{va} and ℓ_{av} . Since the terms ℓ_{va} and ℓ_{av} and the definition of their respective positive and negative sets is symmetric, we will restrict our illustration to the case of ℓ_{va} . Given a video feature vector ν_i^r , we set $\mathcal{P}_i^{va} = \{(\alpha_i^r, 1)\} \cup \text{NN}_k(\alpha_i^r, Q_a) \times W_{1:k}(\alpha_i^r)$, where α_i^r is the feature of the corresponding audio clip with *identical temporal augmentation* (note the superscript). This is in contrast to the definition of ℓ_{vv} where positive pairs were not temporally aligned. As we will show in ablations, we found temporal alignment to be important for cross-modal contrastive learning. The set of negatives is defined as $\mathcal{N}_i^{va} = \{\nu_j \in \mathcal{B} | j \neq i\} \cup \{\alpha_j \in \mathcal{B} | j \neq i\} \cup \text{NN}_{q:q+m}(\alpha_i^r, Q_a)$, *i.e.*, we consider both other audio and other video feature vectors as negatives.

Multi-Modal Contrastive Objective. Our final contrastive objective is composed of the following intra- and inter-modal terms

$$\mathcal{L}_{\text{CRL}} = \mathbb{E}_{\nu_i, \alpha_i} [\ell_{vv}(\nu_i^r, \mathcal{P}_i^{vv}, \mathcal{N}_i^{vv}) + \ell_{va}(\nu_i^r, \mathcal{P}_i^{va}, \mathcal{N}_i^{va}) + \ell_{av}(\alpha_i^r, \mathcal{P}_i^{av}, \mathcal{N}_i^{av})]. \quad (4)$$

Note that our final model does not contain an audio-audio contrastive term. Indeed, we find that including such a term analogous to ℓ_{vv} hurts the final feature performance in transfer experiments (see ablations in Table 3). An illustration of the intra- and inter-modal terms is given in Figure 1.

Temporal Self-Supervision for Video and Audio

Aside from learning from the correspondence between audio and video as proposed above, we also want to promote the learning of temporal features in both domains through self-supervised temporal reasoning tasks. These temporal pretext tasks can be categorized into unitary intra-modal tasks and pairwise intra- and cross-modal objectives (see Figure 2).

Intra-Modal Speed and Direction Classification. To capture short-term temporal video features we leverage the classification of temporal transformations as SSL objectives (Jenni, Meishvili, and Favaro 2020). Concretely, we train the model to predict whether videos are played forward or backward and at which playback speed. The direction classification is a simple binary classification task per clip, and either direction is equally likely during training. The speed

classification is posed as a classification task among 4 speed classes (1 \times , 2 \times , 4 \times , and 8 \times speedup). The speed manipulations are implemented via temporal subsampling, and all the speed classes are equally likely during pre-training.

We propose to leverage such temporal supervision in the audio domain in this work. We apply the temporal transformations to the 1D raw audio signal (analogous to the video domain), *i.e.*, we subsample the signal for speed manipulations and reverse its direction before computing the spectrogram. In experiments, we also investigate an alternative approach where we perform the temporal transformations in the audio spectrogram (thus not manipulating the frequency). Interestingly, we found that transforming the raw audio waveform is much more effective, even when accounting for processing artifacts in manipulating the spectrogram (see ablations in Table 2).

Intra- and Inter-Modal Temporal Ordering. To capture the longer-term dynamics of videos we propose to also perform temporal learning tasks at the clip level by predicting the order of two video clips. Besides performing such temporal ordering solely on video (Jenni and Jin 2021; Xu et al. 2019; Kim, Cho, and Kweon 2019), we extend it to temporal ordering of the audio tracks and cross-modal audio-visual temporal ordering. Concretely, we pose the three-way classification of two temporal signals into 1. correctly ordered, 2. overlapping, and 3. wrongly ordered. This task is implemented by concatenating the representations of the two time-signals along the channel dimension and feeding it through a classifier, *e.g.*, $\phi_{va}([F_v(v_i), F_a(a_i)])$ for video-audio ordering. Likewise, we introduce classifiers ϕ_{vv} , ϕ_{av} , and ϕ_{aa} for video-video, audio-video, and audio-audio temporal ordering.

Finally, we jointly optimize the network weights of the audio and video branch on the combination of the temporal and contrastive objectives. Concretely, let $\mathcal{L}_{\text{TEMP}} = \mathcal{L}_{\text{speed}} + \mathcal{L}_{\text{direction}} + \mathcal{L}_{\text{order}}$ be the sum of all the losses for the above temporal reasoning tasks. The final objective is then given by

$$\mathcal{L}_{\text{SSL}} = \mathcal{L}_{\text{CRL}} + \lambda \mathcal{L}_{\text{TEMP}}, \quad (5)$$

where we set $\lambda = 0.5$.

Implementation Details

For our video encoder F_v we consider variants of the popular 3D-ConvNet architectures R3D (Hara, Kataoka, and Satoh 2018) and R(2+1)D (Tran et al. 2018). If not specified otherwise, input video clips are assumed to contain 16 frames of resolution 112×112 for R(2+1)D, 128×128 for R3D-18, and 224×224 for R3D-34. Our audio encoder F_a is based on a standard ResNet-34 (He et al. 2016) architecture in all experiments. Input spectrograms to the audio encoder are resized to 224×224 .

We train the models using the AdamW optimizer (Loshchilov and Hutter 2017) with a weight decay set to 10^{-4} . The learning rate follows a cosine annealing schedule (Loshchilov and Hutter 2016) with a maximum learning rate of $3 \cdot 10^{-4}$ and linear warm-up in the first training epoch. By default, we train all the models with a batch size of 256.

Besides the temporal input transformations described above (*i.e.*, playback speed+direction changes and temporal cropping), we use the typical data augmentation recipe for contrastive methods, *i.e.*, horizontal flipping, color-jittering, and random spatial cropping. We do not apply any augmentations beyond the temporal ones for audio.

The projection MLPs ψ contain two hidden layers of size 1024 and output feature embeddings of size 256. The prediction MLPs ϕ contain a single hidden layer with a hidden dimension of 1024. We apply synchronized batch norm in both MLPs (including the output of ψ) following prior work (Chen et al. 2020b). The classification heads for the temporal self-supervision tasks follow a similar design to ψ , except that no batch norm is applied to the output in this case.

To evaluate models in transfer experiments, we average predictions of multiple temporal and spatial crops. Likewise, the features for linear probes and nearest-neighbor retrieval are obtained by averaging multiple crops and standardizing the resulting features using the training set statistics.

Experiments

Datasets. As a pre-training dataset we use Kinetics (Zisserman et al. 2017) in most of our experiments. The dataset contains around 350K training videos categorized into 600 human action classes. For transfer experiments we consider UCF101 (Soomro, Zamir, and Shah 2012) and HMDB51 (Kuehne et al. 2011) which are significantly smaller datasets with human action annotations. We use these datasets to evaluate the transfer performance of the video branch, both via fine-tuning to action recognition and as fixed feature extractors for video retrieval. We evaluate the audio branch of our model on ESC50 (Piczak 2015) in terms of environmental audio classification.

Augmented VGG-Sound. Finally, we use the test set of VGG-Sound (Chen et al. 2020a) to evaluate both branches in terms of their robustness to heavy content manipulation for fingerprinting applications. Concretely we generate the following four augmented versions of the dataset by applying different types of audio and video transformations (examples in parenthesis):

1. **AugVGG-IP** - "In-Place" manipulations (V: noise, blur, pixelization, emoji overlay; A: noise, clicks).
2. **AugVGG-S** - "Spatial" transformations (V: cropping, padding, rotation; A: pitch shift, reverb, freq. filter).
3. **AugVGG-T** - "Time" transforms (V+A: speed, crops).
4. **AugVGG-C** - "Combined" (one of each type above).

We use the AugLy library for the dataset creation (Papakipos and Bitton 2022). For fingerprinting evaluations, we report recall at k for these datasets where queries stem from AugVGG-x and retrievals are computed on the clean test set.

Ablations

We perform extensive ablation experiments to investigate the influence of the contrastive loss function design, the various temporal self-supervision signals for audio representation learning, and our combined audio-visual model.

On the Design of the Contrastive Loss. We perform experiments with different variants of the general contrastive ob-

Table 1: **Contrastive Loss Design.** We explore different configurations of the contrastive loss formulation in Eq. 1 in combination with temporal SSL when applied to video-video learning (no audio is being used). We report nearest-neighbor classifier accuracy on UCF101 and HMDB51 and recall @1 for robust video fingerprinting on VGG-Sound.

| Experiment | UCF101 | HMDB51 | AugVGG-C |
|-------------------------|-------------|-------------|-------------|
| | 1-NN | 1-NN | R@1 |
| (a) w/o Q_v positives | 61.5 | 32.5 | 65.5 |
| (b) w/o Q_v negatives | <u>63.9</u> | 34.0 | <u>65.1</u> |
| (c) hard negatives | 63.5 | 33.3 | 65.5 |
| (d) easy negatives | 62.9 | <u>34.8</u> | <u>65.1</u> |
| (e) uniform ω_j | 63.7 | 33.1 | 64.1 |
| Baseline | 65.3 | 35.3 | 65.5 |
| (f) NNCLR | 64.8 | 34.2 | 62.8 |
| (g) SimCLR | 53.9 | 29.2 | 61.1 |
| (h) SimSiam | 62.8 | 34.0 | 60.9 |

jective in Equation 1 and compare it to some popular existing baselines. For faster experimentation, we perform these experiments on video only (we do not use the audio channel here) and pre-train the networks for 40 epochs. We use an R3D-18 network architecture and perform the temporal reasoning tasks among video clips in the experiments. We compare the following variants and report results in Table 1:

(a)-(b) Positives and negatives from the memory bank. In this case, we remove the nearest neighbors from the memory bank as additional positives (a) or remove the negative sampling from Q_v (b). We observe that both positives and negatives from Q_v demonstrate clear benefits, while the positives provide more significant improvements, especially in action retrieval performance.

(d)-(e) Difficulty of negatives. Instead of sampling negatives starting from the median of nearest neighbors in the memory bank, we start at the 90th percentile for hard negatives in (c) and at the 20th percentile for easy negatives (d). Both variants lead to inferior action retrieval performance, and easy negatives hurt fingerprinting.

(f) Equal weighting of positives. Instead of the cross-view similarity-based weighting of the positives, all five positive examples contribute equally to the loss in this case. We observe a drop in the fingerprinting retrieval especially, possibly due to decreased importance of the exact match in the loss. This case is similar to the approach in (Koochpayegani, Tejankar, and Pirsiavash 2021).

(g)-(i) Prior approaches. We replace our proposed loss with existing prior approaches. NNCLR (Dwibedi et al. 2021) replaces the embedding of one view with its nearest neighbor in the memory bank. While this leads to good performance in action retrieval, the performance for fingerprinting suffers. We hypothesize that the lack of the exact match and the lack of additional negatives are the main reason. Key differences to SimCLR (Chen et al. 2020b) are 1. lack of nearest neighbors, 2. lack of predictor MLP, 3. gradient back-propagation through both views. SimCLR requires much larger mini-batches to perform well, which is prohibitive on

Table 2: **Temporal Self-Supervision for Audio Feature Learning.** We explore how the different temporal self-supervision signals impact the audio representation performance for downstream audio classification on ESC50 and audio fingerprinting on VGG-Sound. The audio encoder is pre-trained with temporal supervision and audio-audio contrastive learning (no RGB frames were used).

| Ablation | ESC50 | | AugVGG-C |
|-----------------------|-------------|-------------|-------------|
| | Linear | 1-NN | R@1 |
| (a) w/o speed | 80.4 | 58.4 | 21.1 |
| (b) w/o direction | 79.0 | 56.7 | <u>21.8</u> |
| (c) w/o order | <u>80.6</u> | <u>58.5</u> | 21.5 |
| (d) spect.-resize | 71.0 | 50.8 | 19.3 |
| (e) + rand. STFT-step | 76.5 | 53.3 | 21.5 |
| Baseline | 82.2 | 61.0 | 21.9 |

video. Finally, SimSiam (Chen and He 2020) lacks any negative examples but is otherwise identical to (a). We can again observe the importance of explicit negatives for the fingerprinting use case.

The Benefits of Temporal Self-Supervision for Audio. We performed ablation experiments to demonstrate the different temporal learning tasks’ effect on audio feature performance. We only train the audio branch in these experiments and combine the temporal tasks with an audio-audio contrastive term. Networks were again trained for 40 epochs on Kinetics. In Table 2 (a)-(c), we report the performance of models where each of the three temporal supervision signals is removed. We can observe that each task significantly benefits feature performance, especially in downstream audio recognition tasks. In ablation (d)-(e), the temporal speed transformations are realized by resizing the audio spectrogram instead of subsampling the raw audio signal. We observe clear performance degradations in these cases, even when randomizing the frame step of the STFT, which could prevent some possible shortcuts due to resizing artifacts.

Combined Contrastive and Temporal Audio-Visual Learning. Finally, we validate our combined audio-visual model through experiments demonstrating the importance of the inclusion (or exclusion) of the different contrastive and temporal objectives and ablate model design variations. In this set of experiments, we use an R(2+1)D-18 architecture for the video encoder, and we again train the model for 40 epochs. Table 3 shows the results of the following experiments:

(a)-(d) Training Objectives: We show the influence of the different contrastive intra- and inter-modal objectives in (a)-(c) and the addition of the temporal reasoning tasks in (d). We observe that the cross-modal term brings the most benefit, followed by including the intra-video term. Interestingly, the exclusion of the intra-audio term performs better in all cases. Finally, note how adding temporal self-supervision to the contrastive objectives provides significant gains across the board.

(e)-(f) Implementation Details: We further illustrate the importance of using temporally aligned positives in the cross-modal contrastive term in (e). We believe that the

Table 3: **Audio-Visual Model Ablations.** We perform ablation experiments to demonstrate the influence of the different self-supervised learning signals in our approach (first block) and various implementation details (second block). The video encoder is evaluated in transfer to action recognition on UCF101 and HMDB51, and the audio encoder for classification on ESC50. The fused audio-video feature is used for fingerprinting on VGG-Sound.

| Experiment | UCF101 | HMDB51 | ESC50 | AugVGG-C |
|-------------------|-------------|-------------|-------------|-------------|
| | 1-NN | 1-NN | 1-NN | R@1 |
| (a) w/o A-V CLR | 61.0 | 32.2 | 62.9 | 73.8 |
| (b) w/o V-V CLR | 61.0 | 33.8 | 67.3 | 69.6 |
| (c) w/ A-A CLR | <u>69.1</u> | 37.4 | 68.3 | 78.1 |
| (d) w/o temp.-SSL | <u>67.5</u> | 37.4 | 67.4 | 78.6 |
| (e) unaligned A-V | 68.4 | 37.2 | 68.9 | 78.8 |
| (f) shared Q | 68.8 | <u>38.5</u> | 69.4 | 79.3 |
| Baseline | 70.7 | 40.5 | <u>69.0</u> | 78.1 |

model can leverage the temporal audio-visual correspondence to better associate scene events with their sounds. Finally, in (f), we use only a single memory bank which we feed with the averages of the features from both modalities. Interestingly, this outperforms separate memory banks for fingerprinting and audio recognition.

Comparison to Prior Work on Video SSL

We compare against prior self-supervised video representation learning methods in transfer learning experiments for action recognition and retrieval on UCF101 and HMDB51. We train and evaluate two different video encoders in these comparisons: 1. a smaller-scale experiment with an R(2+1)D-18 trained at 112×112 and 2. a larger-scale experiment with an R3D-34 trained at 224×224 resolution.

Transfer to Action Recognition and Audio Classification. We compare on UCF101 and HMDB51 action recognition and ESC50 audio classification in Table 4, both with full fine-tuning and linear probes when available. A fair comparison to and among prior works is difficult due to significant differences in pre-training datasets, network architectures, input configurations, and training duration. We indicate some of these factors that are known to impact performance in the table. While there are prior works (Recasens et al. 2021; Qian et al. 2020) reporting comparable performance in some tasks, they either use larger architectures, larger pre-training datasets, train for longer, or a combination of those. Our method is more efficient in comparison while still achieving state-of-the-art performance. Notably, when comparing the most common setting using R(2+1)D-18 trained on Kinetics-400, we outperform the best prior results by +3.1%, +9.0%, and +5.7% on UCF101, HMDB51, and ESC-50 respectively.

Video Retrieval Performance. We compare to the prior state-of-the-art approaches TCLR (Dave et al. 2021), GDT (Patrick et al. 2020), Robust-xID (Morgado, Misra, and Vasconcelos 2021), and TE-CVRL (Jenni and Jin 2021) in video retrieval benchmarks on UCF101 and HMDB51 in Table 5.

Table 4: **Action Recognition on UCF101 and HMDB51 and Audio Classification on ESC50.** We report action recognition accuracy after full fine-tuning and linear probe evaluation. We indicate the pre-training dataset, resolution, the number of frames, iterations (or epochs in brackets), and pre-training data modalities (V=RGB, A=audio).

| Method | Dataset | Res. | Frames | It. [Ep.] | Network | Mod. | UCF101 | | HMDB51 | | ESC50 |
|---|---------|------|--------|------------|------------|------|-------------|-------------|-------------|-------------|-------------|
| | | | | | | | FT | Lin. | FT | Lin. | Lin. |
| TE-CVRL (Jenni and Jin 2021) | K400 | 112 | 16 | [200] | R(2+1)D-18 | V | 88.2 | | 62.2 | | |
| CVRL (Qian et al. 2020) | K600 | 224 | 32 | [800] | R3D-50 | V | <u>93.4</u> | <u>90.6</u> | 68.0 | 59.7 | |
| MMV (Alayrac et al. 2020) | AS | 224 | 32 | 500K | R(2+1)D-18 | V+A | 91.5 | 83.9 | 70.1 | 60.0 | |
| BraVe (Recasens et al. 2021) | AS | 224 | 32 | 620K | R(2+1)D-18 | V+A | 93.6 | 90.0 | 70.8 | <u>63.6</u> | |
| AVTS (Korbar, Tran, and Torresani 2018) | K400 | 224 | 25 | [90] | MC3 | V+A | 85.8 | | 56.9 | | 76.7 |
| XDC (Alwassel et al. 2019) | K400 | 224 | 32 | 900K | R(2+1)D-18 | V+A | 84.2 | | 47.1 | | 78.5 |
| GDT (Patrick et al. 2020) | K400 | 112 | 32 | [200] | R(2+1)D-18 | V+A | 88.7 | | 57.8 | | 78.6 |
| AVID (Morgado, Vasconcelos, and Misra 2021) | K400 | 224 | 32 | [400] | R(2+1)D-18 | V+A | 87.5 | | 60.8 | | 79.1 |
| Ours | VGG-S | 112 | 16 | 160K [240] | R(2+1)D-18 | V+A | 90.9 | 86.8 | 70.2 | 55.9 | 87.9 |
| Ours | K400 | 112 | 16 | 200K [240] | R(2+1)D-18 | V+A | 91.8 | 88.0 | 71.2 | 58.2 | 84.8 |
| Ours | K600 | 112 | 16 | 200K [150] | R(2+1)D-18 | V+A | 92.2 | 90.3 | <u>72.2</u> | 62.6 | <u>86.4</u> |
| Ours | K600 | 224 | 16 | 400K [300] | R3D-34 | V+A | 93.6 | 91.8 | 74.6 | 65.8 | 85.5 |




Figure 3: **Video Fingerprinting Performance.** We report instance retrieval performance under video content manipulation on the different AugVGG variants. We show results using a video only (V), audio only (A), and a joint audio-visual model (A+V).

Table 5: **Video Retrieval on UCF101 and HMDB51.** We report recall at k ($R@k$) for k -NN video retrieval. All methods use a R(2+1)D-18 network.

| Method | UCF101 | | | HMDB51 | | |
|--------------------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | R@1 | R@5 | R@20 | R@1 | R@5 | R@20 |
| TCLR | 56.9 | 72.2 | 84.6 | 24.1 | 45.8 | 75.3 |
| GDT | 57.4 | 73.4 | 88.1 | 25.4 | 51.4 | 75.0 |
| Robust-xID | 60.9 | 79.4 | 90.8 | 30.8 | 55.8 | 79.7 |
| TE-CVRL | <u>64.2</u> | <u>81.1</u> | <u>92.6</u> | <u>33.1</u> | <u>60.8</u> | <u>84.1</u> |
| Ours (R(2+1)D-18) | 80.6 | 90.4 | 96.4 | 44.9 | 70.4 | 87.6 |
| Ours (R3D-34) | 85.2 | 93.0 | 97.3 | 51.3 | 74.3 | 91.4 |

Queries stem from the test set, and retrievals are computed on the training set of the respective dataset. A retrieval is assumed correct when the class of query and retrieval agree. We report recall at k for different nearest neighbors. Our model outperforms prior methods by a considerable margin.

Video Fingerprinting Performance on AugVGG. Finally, we report video retrieval performance under video manipulations in Figure 3. We report recall at k for all four datasets and three models: 1. fused audio and video features, 2. video-only, and 3. audio-only. The fused embedding (concatenation of audio and video features) performs best in all cases, followed by the video model. Surprisingly, AugVGG-IP with in-place augmentations is most difficult, while performance on AugVGG-S and AugVGG-T is close to perfect.

Audio-Visual Feature Fusion. We explore the fusion of the aural and visual features learned through our approach for

Table 6: **Modality Fusion.** We explore the fusion of our audio-visual features for downstream video classification.

| Modalities | VGG-Sound | K600 |
|-------------|-------------|-------------|
| Audio | 39.1 | 15.7 |
| Video | <u>39.7</u> | <u>56.8</u> |
| Audio+Video | 53.9 | 58.4 |

downstream video understanding tasks. We compare linear probe accuracy for audio, video, and fused features learned on VGG-Sound and Kinetics-600 in Table 6. Interestingly, combining both modalities improves not only the audio-focused VGG-Sound benchmark but also the appearance-focused classification task on Kinetics-600.

Conclusions

We introduced a novel method to learn video and audio representations by exploiting temporal and audio-visual self-supervision. To learn temporal features, our model learns through time-related pretext tasks, which we extend to the audio domain and the cross-modal setting. We propose a novel contrastive loss design and a model with both intra- and cross-modal contrastive objectives to learn from the audio-visual correspondence in videos. Experiments demonstrate that representations that integrate both temporal and aural features achieve state-of-the-art video classification and retrieval performance.

References

- Akbari, H.; Yuan, L.; Qian, R.; Chuang, W.-H.; Chang, S.-F.; Cui, Y.; and Gong, B. 2021. Vatt: Transformers for multi-modal self-supervised learning from raw video, audio and text. *Advances in Neural Information Processing Systems*, 34.
- Alayrac, J.-B.; Recasens, A.; Schneider, R.; Arandjelović, R.; Ramapuram, J.; De Fauw, J.; Smaira, L.; Dieleman, S.; and Zisserman, A. 2020. Self-supervised multimodal versatile networks. *Advances in Neural Information Processing Systems*, 33: 25–37.
- Alwassel, H.; Mahajan, D.; Korbar, B.; Torresani, L.; Ghanem, B.; and Tran, D. 2019. Self-supervised learning by cross-modal audio-video clustering. *arXiv preprint arXiv:1911.12667*.
- Alwassel, H.; Mahajan, D.; Korbar, B.; Torresani, L.; Ghanem, B.; and Tran, D. 2020. Self-supervised learning by cross-modal audio-video clustering. *Advances in Neural Information Processing Systems*, 33: 9758–9770.
- Arandjelovic, R.; and Zisserman, A. 2017. Look, listen and learn. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 609–617. IEEE.
- Bai, Y.; Fan, H.; Misra, I.; Venkatesh, G.; Lu, Y.; Zhou, Y.; Yu, Q.; Chandra, V.; and Yuille, A. 2020. Can Temporal Information Help with Contrastive Self-Supervised Learning? *arXiv preprint arXiv:2011.13046*.
- Benaim, S.; Ephrat, A.; Lang, O.; Mosseri, I.; Freeman, W. T.; Rubinstein, M.; Irani, M.; and Dekel, T. 2020. Speed-Net: Learning the Speediness in Videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9922–9931.
- Black, A.; Bui, T.; Jenni, S.; Swaminathan, V.; and Colloso, J. 2021. VPN: Video Provenance Network for Robust Content Attribution. In *European Conference on Visual Media Production*, 1–10.
- Brattoli, B.; Büchler, U.; Wahl, A.-S.; Schwab, M. E.; and Ommer, B. 2017. Lstm self-supervision for detailed behavior analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882*.
- Chen, H.; Xie, W.; Vedaldi, A.; and Zisserman, A. 2020a. Vggsound: A large-scale audio-visual dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 721–725. IEEE.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020b. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Chen, X.; and He, K. 2020. Exploring Simple Siamese Representation Learning. *arXiv preprint arXiv:2011.10566*.
- Dave, I.; Gupta, R.; Rizve, M. N.; and Shah, M. 2021. TCLR: Temporal Contrastive Learning for Video Representation. *arXiv preprint arXiv:2101.07974*.
- Doersch, C.; Gupta, A.; and Efros, A. A. 2015. Unsupervised Visual Representation Learning by Context Prediction. *ICCV*.
- Dosovitskiy, A.; Fischer, P.; Springenberg, J. T.; Riedmiller, M.; and Brox, T. 2015. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(9): 1734–1747.
- Dwibedi, D.; Aytar, Y.; Tompson, J.; Sermanet, P.; and Zisserman, A. 2021. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9588–9597.
- Epstein, D.; Chen, B.; and Vondrick, C. 2020. Oops! Predicting Unintentional Action in Video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 919–929.
- Feichtenhofer, C.; Fan, H.; Xiong, B.; Girshick, R.; and He, K. 2021. A large-scale study on unsupervised spatiotemporal representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3299–3309.
- Fernando, B.; Bilen, H.; Gavves, E.; and Gould, S. 2017. Self-supervised video representation learning with odd-one-out networks. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, 5729–5738. IEEE.
- Gidaris, S.; Singh, P.; and Komodakis, N. 2018. Unsupervised Representation Learning by Predicting Image Rotations. In *International Conference on Learning Representations*.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P. H.; Buchatskaya, E.; Doersch, C.; Pires, B. A.; Guo, Z. D.; Azar, M. G.; et al. 2020. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*.
- Hara, K.; Kataoka, H.; and Satoh, Y. 2018. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 6546–6555.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9729–9738.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Jenni, S.; and Favaro, P. 2018. Self-Supervised Feature Learning by Learning to Spot Artifacts. In *CVPR*.
- Jenni, S.; and Jin, H. 2021. Time-equivariant contrastive video representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9970–9980.
- Jenni, S.; Jin, H.; and Favaro, P. 2020. Steering Self-Supervised Feature Learning Beyond Local Pixel Statistics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6408–6417.

- Jenni, S.; Meishvili, G.; and Favaro, P. 2020. Video representation learning by recognizing temporal transformations. *arXiv preprint arXiv:2007.10730*.
- Kim, D.; Cho, D.; and Kweon, I. S. 2019. Self-supervised video representation learning with space-time cubic puzzles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 8545–8552.
- Koohpayegani, S. A.; Tejankar, A.; and Pirsiavash, H. 2021. Mean shift for self-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10326–10335.
- Korbar, B.; Tran, D.; and Torresani, L. 2018. Cooperative learning of audio and video models from self-supervised synchronization. In *Advances in Neural Information Processing Systems*, 7763–7774.
- Kuehne, H.; Jhuang, H.; Garrote, E.; Poggio, T.; and Serre, T. 2011. HMDB: a large video database for human motion recognition. In *Proceedings of the International Conference on Computer Vision (ICCV)*.
- Lee, H.-Y.; Huang, J.-B.; Singh, M.; and Yang, M.-H. 2017. Unsupervised representation learning by sorting sequences. In *Proceedings of the IEEE International Conference on Computer Vision*, 667–676.
- Lee, S.; and Yoo, C. D. 2008. Robust video fingerprinting for content-based video identification. *IEEE Transactions on Circuits and Systems for Video Technology*, 18(7): 983–988.
- Loshchilov, I.; and Hutter, F. 2016. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*.
- Loshchilov, I.; and Hutter, F. 2017. Fixing weight decay regularization in adam. *arXiv preprint arXiv:1711.05101*.
- Misra, I.; Zitnick, C. L.; and Hebert, M. 2016. Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision*, 527–544. Springer.
- Morgado, P.; Misra, I.; and Vasconcelos, N. 2021. Robust audio-visual instance discrimination. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12934–12945.
- Morgado, P.; Vasconcelos, N.; and Misra, I. 2021. Audio-Visual Instance Discrimination with Cross-Modal Agreement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12475–12486.
- Noroozi, M.; and Favaro, P. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, 69–84. Springer.
- Owens, A.; and Efros, A. A. 2018. Audio-Visual Scene Analysis with Self-Supervised Multisensory Features. In *The European Conference on Computer Vision (ECCV)*.
- Owens, A.; Wu, J.; McDermott, J. H.; Freeman, W. T.; and Torralba, A. 2016. Ambient sound provides supervision for visual learning. In *European Conference on Computer Vision*, 801–816. Springer.
- Papakipos, Z.; and Bitton, J. 2022. AugLy: Data Augmentations for Robustness. *arXiv:2201.06494*.
- Patrick, M.; Asano, Y. M.; Fong, R.; Henriques, J. F.; Zweig, G.; and Vedaldi, A. 2020. Multi-modal self-supervision from generalized data transformations. *arXiv preprint arXiv:2003.04298*.
- Piczak, K. J. 2015. ESC: Dataset for environmental sound classification. In *Proceedings of the 23rd ACM international conference on Multimedia*, 1015–1018.
- Qian, R.; Meng, T.; Gong, B.; Yang, M.-H.; Wang, H.; Belongie, S.; and Cui, Y. 2020. Spatiotemporal contrastive video representation learning. *arXiv preprint arXiv:2008.03800*.
- Recasens, A.; Luc, P.; Alayrac, J.-B.; Wang, L.; Strub, F.; Tallec, C.; Malinowski, M.; Pătrăucean, V.; Altché, F.; Valko, M.; et al. 2021. Broaden your views for self-supervised video learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1255–1265.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*.
- Tran, D.; Wang, H.; Torresani, L.; Ray, J.; LeCun, Y.; and Paluri, M. 2018. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 6450–6459.
- Wang, X.; Liu, Z.; and Yu, S. X. 2020. Unsupervised Feature Learning by Cross-Level Discrimination between Instances and Groups. *arXiv preprint arXiv:2008.03813*.
- Wei, D.; Lim, J.; Zisserman, A.; and Freeman, W. T. 2018. Learning and using the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8052–8060.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3733–3742.
- Xu, D.; Xiao, J.; Zhao, Z.; Shao, J.; Xie, D.; and Zhuang, Y. 2019. Self-supervised spatiotemporal learning via video clip order prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 10334–10343.
- Yao, Y.; Liu, C.; Luo, D.; Zhou, Y.; and Ye, Q. 2020. Video Playback Rate Perception for Self-Supervised Spatio-Temporal Representation Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6548–6557.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *European Conference on Computer Vision*, 649–666. Springer.
- Zhang, R.; Isola, P.; and Efros, A. A. 2017. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1058–1067.
- Zisserman, A.; Carreira, J.; Simonyan, K.; Kay, W.; Zhang, B.; Hillier, C.; Vijayanarasimhan, S.; Viola, F.; Green, T.; Back, T.; et al. 2017. The kinetics human action video dataset. *ArXiv*.