

VI-MIX FOR SELF-SUPERVISED VIDEO REPRESENTATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Contrastive representation learning of videos highly rely on exhaustive data augmentation strategies. Therefore, towards designing video augmentation for self-supervised learning, we first analyze the best strategy to mix videos to create a new augmented video sample. Then, the question remains, *can we make use of the other modalities in videos for data mixing?* To this end, we propose **Cross-Modal Manifold Cutmix** (CMMC) that inserts a video tesseract into another video tesseract in the feature space across two different modalities. We find that our video mixing strategy: **Vi-Mix**, i.e. preliminary mixing of videos followed by CMMC across different modalities in a video, improves the quality of learned video representations. We exhaustively conduct experiments for two downstream tasks: action recognition and video retrieval on three popular video datasets UCF101, HMDB51, and NTU-60. We show that the performance of Vi-Mix on both the downstream tasks is on par with the other self-supervised approaches while requiring less training data.

1 INTRODUCTION

The recent advancements in self-supervised representation is credited to the success of using discriminative contrastive loss such as InfoNCE (Gutmann & Hyvärinen, 2010). Given a data sample, contrastive representation learning focus on discriminating its transformed version from a large pool of other instances or their transformations. Thus, the concept of contrastive learning while applicable to any domains, its effectiveness rely on the domain-specific inductive bias as the transformations are obtained from the same data instance. For images, these transformations are usually standard data augmentation techniques (Chen et al., 2020) while in videos, data artifacts that arise from temporal segments within the same video clip (Lee et al., 2017b;a; Fernando et al., 2017; Pickup et al., 2014; Misra et al., 2016).

Recently, data mixing strategies (Zhang et al., 2018; Shen et al., 2020; Yun et al., 2019) have emerged as one of the promising data augmentation for supervised learning methods. These mixing strategies when incorporated with contrastive learning, the quality of the learned representation improves drastically as in Lee et al. (2021); Verma et al. (2021; 2019). Such augmentations introduce semantically meaningful variance for better generalization which is crucial for learning self-supervised representations. While these mixing strategies have been impactful for learning image representations, mixing strategies have been very limitedly explored in the video domain.

Therefore, in this paper, we study the various data mixing strategies for videos, and propose a new approach to overcome their limitation by mixing across modalities. We first investigate and compare the mixing strategies adopted from the image domain, and we find that mixing videos by performing simple interpolation of two video cuboids (Mixup) is more effective than inserting a video cuboid within another (Cutmix). This is in contrast to the observations made in the image domain. Furthermore, unlike learning image representations (Lee et al., 2021), these data mixing strategies are prone to over-fitting when trained for longer, making them limited for videos.

Motivated by the success of previous self-supervised techniques exploiting multiple modalities to learn discriminative video representation as in Arandjelovic & Zisserman (2017a); Chung & Zisserman (2016); Korbar et al. (2018); Arandjelovic & Zisserman (2017b); Owens & Efros (2018); Piergiovanni et al. (2020); Miech et al. (2020), in this paper, we pose the following question: *can we take advantage of other modalities for mixing videos while learning self-supervised representation?*

Different modalities of a video like RGB, optical flow, etc. have different distributions and thus, mixing them directly in the input space makes the task of discriminating similar instances from the other instances easier limiting the quality of the learned representation. To this end, we propose our Cross-Modal Manifold Cutmix (CMMC), that performs data mixing operation ‘across different modalities’ of a video in their hidden intermediate ‘representations’. Given the video encoders from different modalities pre-trained with contrastive loss in addition to mixup augmentation, CMMC exploits the underlying structure of the data manifold. This is done by performing cutmix operation in the feature space across space, time and channels. To the best of our knowledge, this is the first attempt to perform mixing across channels. The channel mixing of the cross-modal feature map enforces the encoder to learn better semantic concepts in the videos. Hence, we train video encoders for different modalities in several stages including the use of mixup strategy in videos and our proposed CMMC. We call this video augmentation strategy - Vi-Mix which stands for Video instance-Mix for contrastive representation learning.

Empirically, we confirm that Vi-Mix being easy to implement, significantly improves contrastive representation learning for videos. We show that Vi-Mix can effectively learn self-supervised representation with small availability of data for pretext task and can also take advantage of other modalities of the videos through manifold mixing strategy. We thoroughly evaluate the quality of the learned representation on two downstream tasks action recognition and retrieval, on UCF101 and HMDB51. We demonstrate the improvement in transferability of the representation learned with Vi-Mix by conducting training on a large scale dataset Kinetics-400 and then finetuning on smaller datasets. Furthermore, we corroborate the robustness of our video data augmentation strategy by observing similar improvements on video skeleton sequences for the task of action recognition.

2 BACKGROUND

In this section, we first review a general contrastive learning mechanism used for learning self-supervised video representation. Then, we review a data mixing formulation for self-supervision in the image domain. Let $\mathcal{X} \in \mathcal{R}^{T \times 3 \times H \times W}$ be a sequence of video. The objective is to learn a mapping $f : \mathcal{X} \rightarrow z$ where $z \in \mathcal{R}^D$, that can be effectively used to discriminate video clips for various downstream tasks, e.g. action recognition, retrieval, etc.

Contrastive Learning. Assume a set of augmentation transformations A is applied to \mathcal{X} . So, for a particular video there exists a positive (say $\tilde{\mathcal{X}}$) whereas the other transformed videos in a mini-batch are considered as negatives. The encoder $f(\cdot)$ and its exponential average model $\tilde{f}(\cdot)$ maps the positives and negatives respectively to embedding vectors. Therefore, the contrastive loss for a sample \mathcal{X}_i is formulated as

$$\mathcal{L}(\mathcal{X}_i) = -\log \frac{\exp(z_i \cdot \tilde{z}_i / \tau)}{\exp(z_i \cdot \tilde{z}_i / \tau) + \sum_{j \in \mathcal{N}} \exp(z_i \cdot z_j / \tau)} \quad (1)$$

where τ is a scaling temperature parameter and \mathcal{N} is the set of negatives. Note that the embedding vectors z_i and \tilde{z}_i are L2-normalized before the loss computation. Thus, the loss \mathcal{L} optimizes the video instances such that the representation of the video instances with the same view are pulled towards each other while pushing away from the other instances.

Data Mix for Contrastive Learning. We revisit the formulation proposed in i-mix (Lee et al., 2021) for mixing data within a batch for contrastive representation learning. Let $y_i \in \{0, 1\}^{BS}$ be the virtual labels of the input \mathcal{X}_i and $\tilde{\mathcal{X}}_i$ in a batch, where $y_{i,i} = 1$ and $y_{i,j \neq i} = 0$. Then, the $(\mathcal{N} + 1)$ -way discrimination loss for a sample in a batch is:

$$\mathcal{L}(\mathcal{X}_i, y_i) = -y_{i,b} \cdot \log \frac{\exp(z_i \cdot \tilde{z}_b / \tau)}{\exp(z_i \cdot \tilde{z}_b / \tau) + \sum_{j \in \mathcal{N}} \exp(z_i \cdot z_j / \tau)} \quad (2)$$

where b ranges from 0 to BS . Thus, the data instances are mixed within a batch for which the loss is defined as:

$$\mathcal{L}_{Mix}((\mathcal{X}_i, y_i), (\mathcal{X}_r, y_r), \lambda) = \mathcal{L}(\text{Mix}(\mathcal{X}_i, \mathcal{X}_r; \lambda), \lambda y_i + (1 - \lambda) y_r) \quad (3)$$

where $\lambda \sim \text{Beta}(\alpha, \alpha)$ is a mixing coefficient, $r \sim \text{rand}(BS)$, and $\text{Mix}()$ is a mixing operator. In the following, we will discuss the appropriate mixing operators in the video domain.

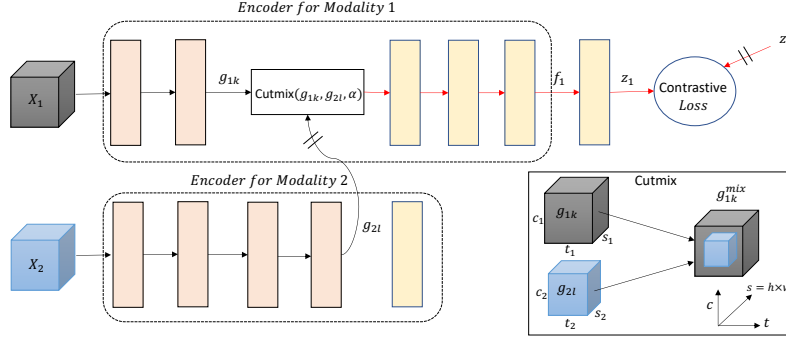


Figure 1: Cross-Modal Manifold Cutmix (CMMC) trains a visual encoder for Modality 1 with \mathcal{X}_1 as input while mixing with input \mathcal{X}_2 from modality 2. The mixing operation between the cross-modal feature maps g_{1k} and g_{2l} is Cutmix which is illustrated at the right.

3 VI-MIX

In this paper, we use the same i-mix formulation (from the above section) for data mixing while learning discriminative self-supervised representation. First, we investigate the best strategies to define the mixing operator for video domain. Furthermore, we introduce a manifold mixing strategy to make use of the other modalities freely available in videos for data mixing. We integrate both these data augmentation strategies, together called Video-instance Mix (Vi-Mix) for contrastive representation learning of Videos.

3.1 MIXING OPERATOR FOR VIDEOS

Unlike mixing operations in images as in Zhang et al. (2018); Verma et al. (2019); Shen et al. (2020); Yun et al. (2019), videos have temporal dimension. We argue that handling temporal dimension in videos is not equivalent to handling spatial dimension in images. For the Mixing operation defined in equation 3, it is straightforward to extend the existing image mixing strategies to videos. Mixup (Verma et al., 2019) in videos perform weighted averaging of two spatio-temporal stack of frames. In contrast to cutmix operator (Yun et al., 2019), mixup operator retains the temporal information in videos and thus facilitates the contrastive representation learning. We empirically corroborate this observation in the experimental analysis. In addition to this, videos possess different modalities like optical flow that can be computed without any supervision. The question remains that can we make use of other modalities in videos for mixing instances while learning contrastive representation? To this end, we introduce **Cross-Modal Manifold Cutmix (CMMC)** strategy for mixing video instances across different modalities which is discussed in the next section.

3.2 CROSS-MODAL MANIFOLD CUTMIX

Different modalities in videos is an additional information that are often exploited for self-supervised learning as in (Han et al., 2020a; Linguo et al., 2021). In contrast to these approaches, we simply propose to mix these different modalities as another data augmentation strategy for self-supervised representation. However, the dissimilarity in distribution between the different modalities (say, RGB and optical flow) in videos makes it harder to mix them at input space. Consequently, we propose Cross-Modal Manifold Cutmix to mix such cross-modal representations in the hidden representation space.

As an extension of the previous notation, we now consider two different modalities \mathcal{X}_{1i} and \mathcal{X}_{2i} for a given video clip \mathcal{X}_i . The objective of the self-supervised task is to learn discriminative video representation, i.e. to learn functions $f_1(\cdot)$ and $f_2(\cdot)$. We decompose the encoder function by $f_1(\mathcal{X}_{1i}) = f_{1k}(g_{1k}(\mathcal{X}_{1i}))$, where g_{1k} is a part of the video encoder for modality 1 with k layers that maps the input data \mathcal{X}_{1i} to a hidden representation. Similarly, f_{1k} maps the hidden representation $g_{1k}(\mathcal{X}_{1i})$ to the embedding vector z_{1i} . Note that we already have trained video encoders $f_{1i}(\cdot)$ and $f_{2i}(\cdot)$ by exploiting the above mentioned Mixup strategy among the video instances in a mini-batch while optimizing the contrastive loss. Now, CMMC is trained in a 4 stage fashion. In the first stage, we train the encoder $f_{1i}(\cdot)$ of modality 1 in 5 steps as illustrated in figure 1. First, we select random

layers k and l from a set of eligible layers in $f_{1i}(\cdot)$ and $f_{2i}(\cdot)$ respectively such that $k \leq l$. This set excludes the input space. Second, we fed a pair of input \mathcal{X}_{1i} and \mathcal{X}_{2r} to their respective video encoders f_1 and f_2 until they reach layer k and layer l respectively. We obtain $g_{1k}(\mathcal{X}_{1i})$ and $g_{2l}(\mathcal{X}_{2r})$ - a hidden representation (spatio-temporal tesseract) of both videos in modality 1 and 2. Third, we perform a data mixing among the hidden representations across two modalities as:

$$g_{1k}^{mix}, \lambda = \text{CutMix}(g_{1k}, g_{2r}; \alpha) \quad (4)$$

$$y_{1k}^{mix} = \lambda y_{1i} + (1 - \lambda) y_{2r} \quad (5)$$

where (y_{1i}, y_{2r}) are one-hot labels, hyper-parameter $\alpha = 1$, and the mixing operator is cutmix as in Yun et al. (2019) which returns the mixing coefficient λ along with the mixed data. For brevity, we omit the input instances in the equation. Fourth, we continue the forward pass in $f_1(\cdot)$ only from layer k to the output embedding, now we denote by z_{1i}^{mix} . Fifth, this embedding is used to compute the $(\mathcal{N} + 1)$ -way discrimination loss which is reformulated as:

$$\mathcal{L}(\mathcal{X}_{1i}, y_{1i}) = -y_{1i,b}^{mix} \cdot \log \frac{\exp(z_{1i}^{mix} \cdot \tilde{z}_{1b}/\tau)}{\exp(z_{1i}^{mix} \cdot \tilde{z}_{1b}/\tau) + \sum_{j \in \mathcal{N}} \exp(z_{1i}^{mix} \cdot z_{1j}/\tau)} \quad (6)$$

The computed gradients are backpropagated through the entire video encoder $f_1(\cdot)$ of modality 1 only. It is to be noted that the video encoder $f_1(\cdot)$ of modality 2 is not trained in this stage. In the second stage, we train the video encoder $f_2(\cdot)$ for modality 2 while freezing the updated learned weights of $f_1(\cdot)$. We continue this cycle twice for each modality, and hence 4 stages to learn the self-supervised video representation in $f_1(\cdot)$ and $f_2(\cdot)$. Algorithm 1 provides the pseudocode of one stage of CMMC for training encoder $f_1(\cdot)$. Thus, to sum up Vi-Mix consists of initially training

Algorithm 1 Pytorch-like style Pseudocode of One stage CMMC for modality 1

```

alpha, mix1 = 1., rand(1, L) # L is the layers in the encoder
mix2 = rand(mix1, L)
x1_q, x1_k = aug(rgb) # Two modalities of the RGB (modality 1)
data
x2 = aug(flow) # Flow data (modality 2)
g1 = f1_q.partial_forward(x1_q, 0, mix1).
g2 = f2_q.partial_forward(x2, 0, mix2)
g_mix, labels_new, lam = CutMix(g1, g2, alpha)
z1 = normalize(f1_q.partial_forward(x1_q, mix1, L))
z2 = normalize(f1_k.forward(x1_k))
z2, g2 = z2.detach(), g2.detach() # no gradient flow
logits = matmul(z1, z2.T) / t
loss = lam * CrossEntropyLoss(logits, arange(len(x1_q)) +
(1 - lam) * CrossEntropyLoss(logits, labels_new)

```

video encoders of modality 1 and modality 2 independently with infoNCE loss as in Chen et al. (2020) and applying mixup augmentation. Then, we perform CMMC among the hidden representations of data from modality 1 and 2 in 4 stages. This is performed by alternation training strategy as in Han et al. (2020a) to make use of the latest learned representations in the cross-modal network. The final learned model is obtained after two cycles of training encoder of each modality.

CutMix in feature space. Here, we explain how the cutmix operator is applied on the video tesseracts in the feature space. Assume that the hidden representation of the input video sequence \mathcal{X}_{1i} in modality 1, $g_{1k} \in \mathcal{R}^{c_1 \times t_1 \times h_1 \times w_1}$, where c_1 represents channel, t_1 time, and $s_1 = h_1 \times w_1$ is the spatial resolution. We generate a new representation g_1^{mix} by combining the hidden representations g_{1k} and g_{2l} . These hidden representations g_{1k} and g_{2l} may differ such that $(c_1, t_1, s_1) \geq (c_2, t_2, s_2)$. Therefore, we define a cutmix operation that combines the video tesseracts in space, time and across channels. We define the combining operation as

$$g_1^{mix} = M \odot g_{1k} + (1 - M) \odot g_{2l} \quad (7)$$

where $M \in \{0, 1\}^{c_1 \times t_1 \times h_1 \times w_1}$ is a binary tensor mask which is decided by sampling the bounding box coordinates $bbox = (b_{c1}, b_{c2}, b_{t1}, b_{t2}, b_{h1}, b_{h2}, b_{w1}, b_{w2})$ from a uniform distribution. In order to preserve the temporal information in a video, we fix $(b_{t1}, b_{t2}) = (0, t_2)$. Similarly, we preserve the channel information processed by the video encoder $f_2(\cdot)$ by fixing $(b_{c1}, b_{c2}) = (0, c_2)$.

Thus, the bounding box selection follows a random sampling of a center coordinate (b_{wc}, b_{hc}) from $(U(0, w_2), U(0, h_2))$. The corner points of the bounding box are determined by

$$b_{w1}, b_{w2} = b_{wc} - \frac{w_2\sqrt{\lambda}}{2}, b_{wc} + \frac{w_2\sqrt{\lambda}}{2} \quad b_{h1}, b_{h2} = b_{hc} - \frac{h_2\sqrt{\lambda}}{2}, b_{hc} + \frac{h_2\sqrt{\lambda}}{2} \quad (8)$$

where $\lambda \sim U(0, 1)$. Even by fixing (b_{t1}, b_{t2}) and (b_{c1}, b_{c2}) , the resultant video tesseract in one modality may not match the dimension of the video tesseracts in other modality across channel and time, if $k < l$. So, we select a 4D bounding box with coordinates $(M_{c1}, M_{c2}, M_{t1}, M_{t2}, M_{h1}, M_{h2}, M_{w1}, M_{w2})$ within the defined binary mask M . We randomly sample a center coordinate $(M_{cc}, M_{tc}, M_{hc}, M_{wc})$ from $(U(0, c_1), U(0, t_1))$, $(U(0, h_1), U(0, w_1))$ respectively. The end points of the binary mask M are determined by

$$\begin{aligned} M_{c1}, M_{c2} &= M_{cc} - \frac{c_2}{2}, M_{cc} + \frac{c_2}{2} & M_{t1}, M_{t2} &= M_{tc} - \frac{t_2}{2}, M_{tc} + \frac{t_2}{2} \\ M_{h1}, M_{h2} &= M_{hc} - \frac{h_2}{2}, M_{hc} + \frac{h_2}{2} & M_{w1}, M_{w2} &= M_{wc} - \frac{w_2}{2}, M_{wc} + \frac{w_2}{2} \end{aligned} \quad (9)$$

For the region within this bounding box, the values in the binary mask is filled with 0, otherwise 1. A new mixing coefficient is computed by $1 - \lambda_{new} = \sum_{c,t,w,h} M_{c,t,w,h}$ denoting the complementary of the proportion of volume occupied by M . This new mixing coefficient λ_{new} is returned by the cutmix function to compute the mixed labels in equation 5.

Thus, we perform a mix operation in videos across all the dimensions including spatial, temporal and channels. However, we preserve the temporal properties of the video instances by retaining a proportion of channel information. This makes cutmix operation effective in the feature space.

4 EXPERIMENTS

In this section, we describe the datasets used in our experimental analysis, implementation details, and evaluation setup. We present ablation studies to illustrate the effectiveness of Vi-Mix video data augmentation and also, provide an exhaustive state-of-the-art comparison with our Vi-Mix models.

4.1 DATASETS

We use two video action recognition datasets: UCF101 (Soomro et al., 2012) and Kinetics-400 (Kay et al., 2017) for self-supervised training of the video encoders. UCF101 contains 13k videos with 101 human actions and Kinetics-400 (K400) contains 240k video clips with 400 human actions. We also use a skeleton action recognition dataset: NTU-RGB+D (Shahroudy et al., 2016) for self-supervised training of a skeleton encoder. NTU-RGB+D (NTU-60) contains 58k videos with 60 human action, all performed indoors. Note that we use the videos or skeleton sequences from the training set only for the self-supervised pre-training. Downstream tasks are evaluated on split1 of UCF101 and HMDB51 (Kuehne et al., 2011), which contains 7k videos with 51 human actions. For evaluation on skeletons, we evaluate on the validation set of NTU-60 on Cross-subject (xsub) and Cross-View (xview) protocols.

4.2 IMPLEMENTATION DETAILS

Vi-Mix is a simple data augmentation strategy that requires cutmix operation in the feature space which is adopted from (Yun et al., 2019) followed by our temporal and channel mixing. The input modalities in our experiments consists of RGB, optical flow and skeletons (3D Poses). The optical flow is computed with the un-supervised TV-L1 algorithm (Sánchez Pérez et al., 2013) and the same pre-processing procedure is used as in Carreira & Zisserman (2017). For the skeleton experiments, the skeleton data $\mathcal{X} \in \mathcal{R}^{C \times T \times V}$ is acquired using KinectV2 sensors, where coordinate feature $C = 3$, # joints $V = 25$, and # frames $T = 50$. Following the pre-processing steps in Linguo et al. (2021), we compute the joints and motion cues. For all the RGB and optical flow models, we choose S3D (Xie et al., 2018) architecture as the backbone whereas for the skeleton model, we choose ST-GCN (Yan et al., 2018) with channels in each layer reduced by 1/4 times as the backbone. For self-supervised representation learning, we adopt a momentum-updated history queue to cache a large number of video features as in MoCo (He et al., 2019). We attach a non-linear projection head, and remove it for downstream task evaluations as done in SimCLR (Chen et al., 2020).

For our experiments with RGB and optical flow, we use $32 \times 128 \times 128$ frames of RGB (or flow) input, at 30 fps. For additional data augmentation, we apply clip-wise consistent random crops, horizontal flips, Gaussian blur and color jittering. We also apply random temporal cropping from the same video as used in Han et al. (2020a). For training a MoCo model with Vi-Mix data augmentation, we initially train the RGB and Flow networks for 300 epochs with mixup data augmentation

Table 1: Different Video mixing strategies (at left) are evaluated on downstream action classification and retrieval tasks. All the models are trained on training samples of UCF101 for 200 epochs with RGB input and tested on the validation set of UCF101 and HMDB51.

Mixing Strategies	Action classification		Retrieval	
	Linear probe		R@1	
	UCF101	HMDB51	UCF101	HMDB51
MoCo (Baseline)	38.2	15.3	29.1	9.8
+ Mixup	49.6	24.9	36.2	16.2
+ Temporal cutmix	41.7	17.3	30.2	12.9
+ ST cutmix	19.8	14.3	32.6	14.3
+ VideoMix	45.5	21.2	34.5	15.2




Figure 2: Accuracy graph illustrating the improvements with CMMC compared to representative base-lines without using mix, manifold mix, and alternate training in CMMC.

independently. The mixup operation is applied in the input space. Then, we train these pre-train networks with CMMC in 4 stages. In each stage, a network with one input modality is trained for 100 epochs by freezing the network with other modality. In the next stage, we reverse the cross-modal networks and continue training the network with other modality. Finally, after the 4 stages, the resultant models are hence trained for 500 epochs in total. For optimization, we use Adam with 10^{-3} learning rate and 10^{-5} weight decay. All the experiments are trained on 4 and 2 V100 GPUs for K400 and others respectively, with a batch size of 32 videos per GPU.

For our experiments with skeleton sequence, we choose Shear with shearing amplitude 0.5 and Crop with a padding ratio of 0.6 as the augmentation strategy as used in Linguo et al. (2021). Note that, for CMMC on skeleton data, we perform cutmix operation only on skeleton vertices followed by channel and temporal mixing. For training with CMMC, the GCN encoders are initially trained for 150 epochs on Joint and Motion cues. This is followed by 2 stage training each with 150 epochs, where encoder with one modality is trained and the other is frozen. For optimization, we use SGD with momentum (0.9) and weight decay (0.0001). The model is trained on 1 V100 with a batch size of 128 skeleton sequences.

4.3 EVALUATION SETUP FOR DOWNSTREAM TASKS

For experiments with RGB and optical flow, we evaluate on two downstream tasks: (i) **action classification** and (ii) **retrieval**. For action classification, we evaluate on (1) **linear probe** where the entire encoder is frozen and a single linear layer followed by a *softmax* layer is trained with cross-entropy loss, and (2) **finetune** where the entire encoder along with a linear and *softmax* layer is trained with cross-entropy loss. Note that the encoders are initialized with the Vi-Mix learned weights. More details for training the downstream action classification framework is provided in the Appendix. For action retrieval, the extracted features from the encoder pre-trained with Vi-Mix are used for nearest-neighbor (NN) retrieval. We report Recall at k ($R@k$) which implies, if the top k nearest neighbours comprise one video pertaining to the same class, a correct retrieval is counted.

4.4 ABLATION STUDIES ON VI-MIX

In this section, we empirically show the correctness of our data augmentation strategy for videos. We also investigate the potential reasons behind the significant improvement of performance with Vi-Mix by conducting relevant experiments.

Which mixing strategy is the best for uni-modal video understanding? In Table 1, we investigate different video mixing strategies based on mixup and cutmix operator for downstream action classification and retrieval tasks. For the augmentations based on Cutmix (Yun et al., 2019), we randomly select a sub-cuboid and plug it into another video. We also consider Videomix (Yun et al., 2020) that performs a cutmix operation across all the frames clipwise consistent. For the virtual labels, we perform label smoothing as defined in equation 3. In image domain, cutmix outperforms the mixup strategy in supervised settings (Yun et al., 2019). However, we find that all strategies using cutmix in temporal dimension (Temporal cutmix), spatio-temporal dimension (ST cutmix) and spatial cutmix (VideoMix) performs worse than simple Mixup strategy. This is because cutmix operation destroys the temporal structure of the videos which is crucial for understanding actions in videos. Similarly, VideoMix where cutmix is performed spatially and not temporally, introduces new contextual information in videos in arbitrary spatial locations. This not only hampers the motion patterns present in the original video but also weakens the similarity between the positive samples in the contrastive loss. Thus, video mixing operators must ensure retention of temporal characteristics in videos.

Table 2: Comparison of different Cross-Modal manifold mixing strategies. Random CV mix layer indicates if the feature map from the other modality is obtained from the same layer as the primary modality or not. $s = h \times w$, t , and c represents spatial, temporal and channel mixing of the feature maps. All the models are trained on training samples of UCF101 for 500 epochs. \mathcal{M} indicates use of cross-modal information.

	Cross-Modal Manifold Mixing strategies	\mathcal{M}	Random CM mix layer	s	t	c	Action classification		Retrieval	
							Linear probe		R@1	
							UCF101	HMDB51	UCF101	HMDB51
RGB	MoCo (Baseline)	×	×	×	×	×	46.8	23.1	33.1	15.2
	+ mixup	×	×	✓	✓	✓	52.8	24.4	37.6	17.6
	+ CM mixup	✓	×	✓	✓	✓	53.9	25.1	40.3	17.6
	+ CM cutmix	✓	×	✓	×	×	54.6	25.5	40.4	17.8
		✓	×	✓	✓	×	55.1	27.1	40.6	16.5
		✓	×	✓	✓	✓	55.2	27.8	41.6	18.6
		✓	✓	✓	✓	✓	55.8	28.3	42.8	19.1
Optical Flow	MoCo (Baseline)	×	×	×	×	×	66.8	30.3	45.2	20.8
	+ mixup	×	×	✓	✓	✓	68.6	33.1	48.7	19.5
	+ CM mixup	✓	×	✓	✓	✓	70.4	33.1	51.2	21.0
	+ CM cutmix	✓	×	✓	×	×	70.4	33.3	51.4	21.0
		✓	×	✓	✓	×	71.8	34.7	52.7	23.1
		✓	×	✓	✓	✓	71.5	33.9	53.1	21.5
		✓	✓	✓	✓	✓	72.4	34.9	53.9	23.1
Two-stream	MoCo (Baseline)	×	×	×	×	×	68.1	33.1	49.8	21.9
	+ mixup	×	×	✓	✓	✓	71.3	36.3	53.8	24.5
	+ CM mixup	✓	×	✓	✓	✓	72.2	35.9	56.1	25.3
	+ CM cutmix	✓	×	✓	×	×	72.1	35.8	56.9	25.1
		✓	×	✓	✓	×	73.1	37.2	56.9	27.3
		✓	×	✓	✓	✓	73.4	37.2	56.4	26.5
		✓	✓	✓	✓	✓	74.0	38.1	58.1	27.1

Why do we need multi-modal mixing strategy for videos? In fig. 2, we illustrate the downstream action classification accuracy vs # epochs plot. This plot clearly shows the importance of applying data mixing augmentation. However, the model trained without multi-modal mixing strategy (CMMC) over-fits after 300 epochs, whereas the models training with multi-modal mixing are still learning discriminative representation. We find that the mixing strategy on video representation learning induces faster training and with CMMC, the models learn cross-modal knowledge without using complicated knowledge distillation techniques as in Crasto et al. (2019); Garcia et al. (2018). It is to be noted that training with cross-modal data augmentation is more beneficial when trained in alternation strategy. In fig. 2, we show that the RGB model with alternate training outperforms the RGB model which is trained for 200 epochs straightaway with the outdated optical flow model. The alternate training strategy takes benefit of the most updated cross-modal model for data mixing and hence learning more discriminative representation.

Diagnosis of CMMC. In Table 2, we provide the results for different configurations of data mixing in the feature space. The objective is to understand the strategies responsible for boosting the performance of the models on UCF101 and HMDB51 for downstream tasks. All the models are initialized with weights obtained from pre-training with mixup for 300 epochs. First, we show that cross-modal mixup (indicated by + mixup) in the feature space exploiting the cross-modal representation outperforms the traditional manifold mixup (indicated by + mixup) in the feature space (Verma et al., 2019) on UCF-101 which does not make use of the cross-modal representation. However, we observe that the action classification accuracy on HMDB51 using optical flow is equivalent for both strategies with or without using cross-modality. This is because HMDB51 mostly consists of static actions with improminent motion patterns which limits the optical flow model to learn motion dominated representation. As a result, this also affects the action classification accuracy on HMDB51 when evaluated with both the streams.

Next, we show the influence of mixing different dimensions in the hidden representation of a video. We perform experiments with cross-modal cutmix occurring in the spatial dimension (s), spatial and temporal (t) dimensions, and finally all the dimensions including channels (c). Note that the manifold cutmix operation is performed within the same data manifold, i.e. the cutmix is performed across the same random layer between RGB and Flow networks. In contradiction to our previous observation of video mixing in the input space, here the temporal cutmix provides a minor boost to the performance of the downstream tasks. This is supported by the fact that hidden representations of a video retains temporal information due to the preceding convolutional operations on the input sample. Also, the channel mixing further boosts the performance by retaining lost temporal information in the resultant mixed feature map. Finally, we introduce more randomness in CMMC by

Table 3: Comparison to the state-of-the-art methods for downstream action classification on UCF101 and HMDB51 for different combination of modalities - video (*V*), audio (*A*), text (*T*) and flow (*F*). For a fair comparison with SOTA, we group the results based on the dataset used for pre-training. 'Frozen ✓' indicates linear probe evaluation and 'Frozen ×' indicates end-to-end finetuning of the encoders.

Method	Network	Res.	Depth	Modality	Dataset	Frozen	UCF101	HMDB51
CoCLR (Han et al., 2020a)	S3D	128	23	V+F	UCF101	✓	72.1	40.2
Vi-Mix	S3D	128	23	V+F	UCF101	✓	74.0	38.1
MemDPC (Han et al., 2020b)	R-2D3D	224	33	V+F	K400	✓	54.1	30.5
VideoMoCo (Pan et al., 2021)	R(2+1)D	112	26	V	K400	✓	78.7	49.2
CoCLR (Han et al., 2020a)	S3D	128	23	V+F	K400	✓	77.8	52.4
CVRL (Qian et al., 2020)	R3D	224	49	V	K400	✓	89.8	58.3
Vi-Mix	S3D	128	23	V+F	K400	✓	79.1	50.4
OPN (Lee et al., 2017a)	VGG	227	14	V	UCF101	×	59.6	23.8
VCOP (Xu et al., 2019)	R(2+1)	112	26	V	UCF101	×	72.4	30.9
CoCLR (Han et al., 2020a)	S3D	128	23	V+F	UCF101	×	87.3	58.7
Vi-Mix	S3D	128	23	V+F	UCF101	×	87.5	59.1
XDC (Alwassel et al., 2020)	R(2+1)D	224	26	V+A	K400	×	84.2	47.1
AVTS (Korbar et al., 2018)	I3D	224	22	V+A	K400	×	83.7	53.0
MemDPC (Han et al., 2020b)	R-2D3D	224	33	V+F	K400	×	86.1	54.5
GDT (Patrick et al., 2021)	R(2+1)D	112	26	V+A	K400	×	89.3	60.0
CoCLR (Han et al., 2020a)	S3D	128	23	V+F	K400	×	90.6	62.9
CVRL (Qian et al., 2020)	R3D	224	49	V	K400	×	92.1	65.4
Vi-Mix	S3D	128	23	V+F	K400	×	91.1	62.3
CBT (Sun et al., 2019)	S3D	112	23	V	K600	×	79.5	44.6
MIL-NCE (Miech et al., 2019)	S3D	224	23	V+T	HTM	×	91.3	61.0
ELO (Piergiovanni et al., 2020)	R(2+1)D	224	65	V+A+F	Youtube8M	×	93.8	67.4
BraVe (Recasens et al., 2021)	TSM-50	224	49	V+F	K600	×	93.8	69.7
Supervised (Xie et al., 2018)	S3D	128	23	V	ImNet+K400	×	96.8	75.9

Table 4: Comparison to the state-of-the-art methods on Nearest-Neighbour video retrieval on UCF101 and HMDB51. Testing set clips are used to retrieve training set videos and $R@k$ is reported for $k \in \{1, 5, 10, 20\}$.

Method	Dataset	UCF101				HMB51			
		R@1	R@5	R@10	R@20	R@1	R@5	R@10	R@20
Jigsaw (Noroozi & Favaro, 2016)	UCF	19.7	28.5	33.5	40.0	-	-	-	-
OPN (Lee et al., 2017a)	UCF	19.9	28.7	34.0	40.6	-	-	-	-
RL-method (Büchler et al., 2018)	UCF	25.7	36.2	42.2	49.2	-	-	-	-
VCOP (Xu et al., 2019)	UCF	14.1	30.3	40.4	51.1	7.6	22.9	34.4	48.8
VCP (Luo et al., 2020)	UCF	18.6	33.6	42.5	53.5	7.6	24.4	36.3	53.6
MemDPC (Han et al., 2020b)	UCF	20.2	40.4	52.4	64.7	7.7	25.7	40.6	57.7
SpeedNet (Benaim et al., 2020)	K400	13.0	28.1	37.5	49.5	-	-	-	-
CoCLR (Han et al., 2020a)	UCF	55.9	70.8	76.9	82.5	26.1	45.8	57.9	69.7
Vi-Mix	UCF	58.1	76.5	83.4	88.7	27.1	50.7	65.1	77.1

randomizing the selection of the cross-modal network layer (refer to `mix2` in algorithm 1) where the mixing takes place. This enables CMMC to take advantage of the features from later layers of the cross-modal network.

4.5 COMPARISON TO THE STATE-OF-THE-ART

In this section, we compare Vi-Mix with previous self-supervised approaches for video/skeleton action classification and video action retrieval. In Table 3, we provide the action classification results on UCF101 and HMDB51 for linear-probing and full finetuning of video encoders with models trained on UCF101 and K400. For linear probing, Vi-Mix only with its strong data augmentation (mixup + CMMC) outperforms CoCLR which shares the same evaluation setting with Vi-Mix, by 1.9% on UCF101. Similar, observation is made for models trained with K400. We also note that our finetuned Vi-Mix encoders outperform approaches using higher spatial resolution (XDC, AVTS, MemDPC), deeper layers (XDC, MemDPC, GDT). However, the lower action classification accuracy of Vi-Mix on HMDB51 compared to CoCLR indicates the requirement of positive mining of data samples in contrastive learning as performed in CoCLR. The large performance gap of Vi-Mix with CVRL is owing to the deeper layers of R3D (49 vs 23) and the input spatial resolution (224 vs 128). Interestingly, our Vi-Mix model pre-trained on K400 performs on par with the models trained on larger datasets substantiating the impact of simple cross-modal video augmentation.

In Table 4, we provide the video retrieval results on UCF101 and HMDB51 for the Vi-Mix models trained on UCF101. This is a classical test for verifying if the pre-trained model learns semantic information while learning self-supervised representation. We test if a query instance clip and its nearest neighbours belong to the same category. Our Vi-Mix model outperforms all the representative baselines by a significant margin on both the datasets.

Table 5: Comparison to the state-of-the-art methods on NTU-60 for action classification using different modalities. We indicate the type of pretext task used in the SOTA methods. * indicates that results reproduced on our settings.

Method	Modality	Network	Pretext task	NTU-60	
				xsub	xview
LongTGAN (Zheng et al., 2018)	Joint	GRU	reconstruction	39.1	48.1
MS ² L (Lin et al., 2020)	Joint	GRU	multi-tasks	52.6	-
AS-CAL (Rao et al., 2021)	Joint	LSTM	contrastive	58.5	64.8
P&C (Su et al., 2020)	Joint	GRU	reconstruction	50.7	76.3
SeBiReNet (Nie et al., 2020)	Joint	GRU	reconstruction	-	79.7
2s-CrosSCLR Linguo et al. (2021)	Joint + Motion	ST-GCN	contrastive	74.5	82.1
3s-CrosSCLR Linguo et al. (2021)	Joint + Motion + Bone	ST-GCN	contrastive	77.8	83.4
SkeletonCLR* (Baseline)	Joint + Motion	ST-GCN	contrastive	70.1	77.2
SkeletonCLR + CMMC	Joint + Motion	ST-GCN	contrastive	72.5	79.1

In Table 5, we generalize our Vi-Mix strategy for skeleton action representation on NTU-60. Since, mixup in the input space hampers the spatial configuration of the skeletons processed by ST-GCN, we only perform CMMC with hidden skeleton representations in the encoder. We treat Joints and Motion as different input modalities of given skeleton data. Note that the cutmix operation across spatial dimension represents the joint vertices (1-dimensional). The downstream action classification results of a skeleton model pre-trained with contrastive learning (SkeletonCLR) using CMMC outperforms its baseline by 2.4% on cross-subject and by 1.% on cross-view protocol. The superior results of CrosSCLR is owing to its cross-modal positive mining which benefits the vanilla SkeletonCLR models. We believe that such positive mining approaches in contrastive learning such as CoCLR or CrosSCLR can benefit by using our video mixing strategies.

5 RELATED WORK

Deep neural networks, especially the networks fabricated for processing videos are data-hungry. While annotating large scale video data is expensive, recently many self-supervised video representation learning approaches have been proposed to make use of the abundant web videos. On one hand, some methods have exploited the temporal structure of the videos, such as predicting if frames appear in order, reverse order, shuffled, color-consistency across frames, etc (Lee et al., 2017b;a; Fernando et al., 2017; Pickup et al., 2014; Misra et al., 2016; Wang et al., 2019; 2017; Vondrick et al., 2018; Recasens et al., 2021). On the other hand, some methods have been taking advantage of the multiple modalities of videos like audio, text, optical flow, etc by designing pretext tasks for their temporal alignment (Chung & Zisserman, 2016; Korbar et al., 2018; Arandjelovic & Zisserman, 2017b; Owens & Efros, 2018; Piergiovanni et al., 2020; Miech et al., 2020; Arandjelovic & Zisserman, 2017a).

Meanwhile, data mixing strategies have gained popularity in image-domain data augmentations for supervised learning (Zhang et al., 2018; Shen et al., 2020; Yun et al., 2019) in addition to their usage also for learning self-supervised image representation (Verma et al., 2019; Lee et al., 2021; Verma et al., 2021). A recent work (unpublished), in the spirit of data mixing in the video domain, VideoMix creates a new training video by inserting a video cuboid into another video in the supervised setting (Yun et al., 2020). In contrast, we focus on mixing video samples for self-supervised representation. Different from the observations in VideoMix, we note that mixup in Vi-Mix is a better augmentation tool rather than strategies involving removal of spatio-temporal sub-space from the original videos. The most closest to our work, Manifold mixup (Verma et al., 2019) focuses on interpolating hidden representation of the samples within a mini-batch, whereas, our proposed CMMC in Vi-Mix performs cutmix operation in the data manifold across different modalities. In addition, we also introduce the notion of channel mixing in the feature space. We find that Vi-Mix is simple to implement while is a strong data augmentation tool for learning self-supervised video representation even with small data size.

6 CONCLUSION

We have analyzed the augmentation strategies for learning self-supervised video representation. We have introduced Vi-Mix which includes performing video mixup followed by Cross-modal manifold mixup to take advantage of additional modalities present in videos. Vi-Mix improves the quality of learned representation and thus brings significant improvement in the performance of downstream tasks on UCF101, HMDB51 and NTU-60 datasets. We believe that Vi-Mix can be a standard video augmentation tool while learning any multi-modal self-supervised video representation.

REFERENCES

- Humam Alwassel, Dhruv Mahajan, Bruno Korbar, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. *CoRR*, abs/1705.08168, 2017a. URL <http://arxiv.org/abs/1705.08168>.
- Relja Arandjelovic and Andrew Zisserman. Objects that sound. *CoRR*, abs/1712.06651, 2017b. URL <http://arxiv.org/abs/1712.06651>.
- Sagie Benaim, Ariel Ephrat, Oran Lang, Inbar Mosseri, William T. Freeman, Michael Rubinstein, Michal Irani, and Tali Dekel. Speednet: Learning the speediness in videos. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Uta Büchler, Biagio Brattoli, and Björn Ommer. Improving spatiotemporal self-supervision by deep reinforcement learning. *CoRR*, abs/1807.11293, 2018. URL <http://arxiv.org/abs/1807.11293>.
- Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4724–4733. IEEE, 2017.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *Workshop on Multi-view Lip-reading, ACCV*, 2016.
- Nieves Crasto, Philippe Weinzaepfel, Karteek Alahari, and Cordelia Schmid. MARS: Motion-Augmented RGB Stream for Action Recognition. In *CVPR*, 2019.
- Basura Fernando, Hakan Bilen, Efstratios Gavves, and Stephen Gould. Self-supervised video representation learning with odd-one-out networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. URL <https://ivi.fnwi.uva.nl/isis/publications/2017/FernandoCVPR2017>.
- Nuno C. Garcia, Pietro Morerio, and Vittorio Murino. Modality distillation with multiple stream networks for action recognition. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (eds.), *Computer Vision – ECCV 2018*, pp. 106–121, Cham, 2018. Springer International Publishing. ISBN 978-3-030-01237-3.
- Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In Yee Whye Teh and Mike Titterton (eds.), *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pp. 297–304, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR. URL <https://proceedings.mlr.press/v9/gutmann10a.html>.
- Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. In *Neurips*, 2020a.
- Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning. In *European Conference on Computer Vision*, 2020b.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019.
- 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.

- Bruno Korbar, Du Tran, and Lorenzo Torresani. Co-training of audio and video representations from self-supervised temporal synchronization. *CoRR*, abs/1807.00230, 2018. URL <http://arxiv.org/abs/1807.00230>.
- Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *2011 International Conference on Computer Vision*, pp. 2556–2563. IEEE, 2011.
- Hsin-Ying Lee, Jia-Bin Huang, Maneesh Kumar Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequence, 2017a.
- Hsin-Ying Lee, Jia-Bin Huang, Maneesh Kumar Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequence, 2017b.
- Kibok Lee, Yian Zhu, Kihyuk Sohn, Chun-Liang Li, Jinwoo Shin, and Honglak Lee. i-mix: A domain-agnostic strategy for contrastive representation learning. In *ICLR*, 2021.
- Lilang Lin, Sijie Song, Wenhan Yang, and Jiaying Liu. Ms2l: Multi-task self-supervised learning for skeleton based action recognition. In *Proceedings of the 28th ACM International Conference on Multimedia*, MM ’20, pp. 2490–2498, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379885. doi: 10.1145/3394171.3413548. URL <https://doi.org/10.1145/3394171.3413548>.
- Li Linguo, Wang Minsi, Ni Bingbing, Wang Hang, Yang Jiancheng, and Zhang Wenjun. 3d human action representation learning via cross-view consistency pursuit. In *CVPR*, 2021.
- Dezhao Luo, Chang Liu, Yu Zhou, Dongbao Yang, Can Ma, Qixiang Ye, and Weiping Wang. Video cloze procedure for self-supervised spatio-temporal learning. *arXiv preprint arXiv:2001.00294*, 2020.
- Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *ICCV*, 2019.
- Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-End Learning of Visual Representations from Uncurated Instructional Videos. In *CVPR*, 2020.
- Ishan Misra, C. Lawrence Zitnick, and Martial Hebert. Shuffle and Learn: Unsupervised Learning using Temporal Order Verification. In *ECCV*, 2016.
- Qiang Nie, Ziwei Liu, and Yunhui Liu. Unsupervised 3d human pose representation with viewpoint and pose disentanglement. In *European Conference on Computer Vision (ECCV)*, 2020.
- Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016.
- Andrew Owens and Alexei A. Efros. Audio-visual scene analysis with self-supervised multisensory features, 2018.
- Tian Pan, Yibing Song, Tianyu Yang, Wenhao Jiang, and Wei Liu. Videomoco: Contrastive video representation learning with temporally adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11205–11214, 2021.
- Mandela Patrick, Yuki M. Asano, Polina Kuznetsova, Ruth Fong, João F. Henriques, Geoffrey Zweig, and Andrea Vedaldi. Multi-modal self-supervision from generalized data transformations, 2021.
- 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 *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2043–2050, 2014. doi: 10.1109/CVPR.2014.262.

- AJ Piergiovanni, Anelia Angelova, and Michael S. Ryoo. Evolving losses for unsupervised video representation learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge J. Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning. *CoRR*, abs/2008.03800, 2020. URL <https://arxiv.org/abs/2008.03800>.
- Haocong Rao, Shihao Xu, Xiping Hu, Jun Cheng, and Bin Hu. Augmented skeleton based contrastive action learning with momentum lstm for unsupervised action recognition. *Information Sciences*, 569:90–109, 2021. doi: <https://doi.org/10.1016/j.ins.2021.04.023>.
- Adrià Recasens, Pauline Luc, Jean-Baptiste Alayrac, Luyu Wang, Florian Strub, Corentin Tallec, Mateusz Malinowski, Viorica Patraucean, Florent Altché, Michal Valko, Jean-Bastien Grill, Aàron van den Oord, and Andrew Zisserman. Broaden your views for self-supervised video learning. *CoRR*, abs/2103.16559, 2021. URL <https://arxiv.org/abs/2103.16559>.
- Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+d: A large scale dataset for 3d human activity analysis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- Dinghan Shen, Mingzhi Zheng, Yelong Shen, Yanru Qu, and Weizhu Chen. A simple but tough-to-beat data augmentation approach for natural language understanding and generation. *arXiv preprint arXiv:2009.13818*, 2020.
- Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, abs/1212.0402, 2012. URL <http://arxiv.org/abs/1212.0402>.
- Kun Su, Xiulong Liu, and Eli Shlizerman. Predict & cluster: Unsupervised skeleton based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9631–9640, 2020.
- Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Contrastive bidirectional transformer for temporal representation learning. *CoRR*, abs/1906.05743, 2019. URL <http://arxiv.org/abs/1906.05743>.
- Javier Sánchez Pérez, Enric Meinhardt-Llopis, and Gabriele Facciolo. TV-L1 Optical Flow Estimation. *Image Processing On Line*, 3:137–150, 2013. doi: 10.5201/ipol.2013.26.
- Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 6438–6447, Long Beach, California, USA, 09–15 Jun 2019. PMLR. URL <http://proceedings.mlr.press/v97/verma19a.html>.
- Vikas Verma, Minh-Thang Luong, Kenji Kawaguchi, Hieu Pham, and Quoc V Le. Towards domain-agnostic contrastive learning. In *International Conference on Machine Learning (ICML)*, 2021.
- Carl Vondrick, Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, and Kevin Murphy. ”tracking emerges by colorizing videos”. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- Xiaolong Wang, Kaiming He, and Abhinav Gupta. Transitive invariance for self-supervised visual representation learning. In *ICCV*, 2017.
- Xiaolong Wang, Allan Jabri, and Alexei A. Efros. Learning correspondence from the cycle-consistency of time. In *CVPR*, 2019.

- Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (eds.), *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XV*, volume 11219 of *Lecture Notes in Computer Science*, pp. 318–335. Springer, 2018. doi: 10.1007/978-3-030-01267-0_19. URL https://doi.org/10.1007/978-3-030-01267-0_19.
- Dejing Xu, Jun Xiao, Zhou Zhao, Jian Shao, Di Xie, and Yueting Zhuang. Self-supervised spatiotemporal learning via video clip order prediction. In *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018.
- Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *International Conference on Computer Vision (ICCV)*, 2019.
- Sangdoon Yun, Seong Joon Oh, Byeongho Heo, Dongyoon Han, and Jinhyung Kim. Videomix: Rethinking data augmentation for video classification. *arXiv preprint arXiv:2012.03457*, 2020.
- Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=r1Ddp1-Rb>.
- Nenggan Zheng, Jun Wen, Risheng Liu, Liangqu Long, Jianhua Dai, and Zhefeng Gong. Unsupervised representation learning with long-term dynamics for skeleton based action recognition. In *AAAI*, pp. 2644–2651, 2018.

A APPENDIX

A.1 MORE IMPLEMENTATION DETAILS

Training/Testing specification for downstream finetuning on UCF101 and HMDB51. At training, we apply the same data augmentation as in the pre-training stage mentioned in section 4.2, except for Gaussain blurring. The model is trained with similar optimization configuration as in the pre-training stage for 500 epochs. At inference, we perform spatially fully convolutional inference on videos by applying ten crops (center crop and 4 corners with horizontal flipping) and temporally take clips with overlapping moving windows. The final prediction is the average *softmax* scores of all the clips.

Training/Testing specification for downstream finetuning on NTU-60. For training the pre-trained ST-GCN along with the linear classifier, we apply the same data augmentation as in the pre-training stage. We train for 100 epochs with learning rate 0.1 (multiplied by 0.1 at epoch 80).

History Queue in MoCo. We adopt momentum-updated history queue as in MoCo (He et al., 2019) to cache a large number of visual features while learning contrastive representation. For our pre-training experiments, we use a *softmax* temperature $\tau = 0.07$, and a momentum $m = 0.999$. The queue size of MoCo for pre-training experiments on UCF101, K400 and NTU-60 are 2048, 16384 and 32768 respectively.

A.2 REGULARIZATION EFFECT OF Vi-MIX

In fig. 3, we provide (1) a plot of training loss of two models, one using Vi-Mix and the other model is using standard data augmentations, and (2) the (K+1)-way accuracy of the pretext task of the models learning contrastive representation. We observe a disparity between the training losses (at left of the figure) in both the models with and without using Vi-Mix. This is owing to the hardness of the pretext task which can be directly correlated with the difficulty of the data transformation, via Vi-Mix data augmentation. Meanwhile, we also note that the (K+1)-way accuracy of the Vi-Mix model




Figure 3: Training on UCF101 dataset. At left, we present the training loss of two models, one using video augmentation: Vi-Mix and the other not. At right, we provide the (K+1)-way accuracy on the pretext task while learning contrastive representation.

while training on contrastive loss is lower than that of the model without using Vi-Mix (at the right of the figure). However, the performance gain of the Vi-Mix model on downstream classification and retrieval tasks shows the regularizing capability of using Vi-Mix type data augmentation.