

Multimodal Contrastive Learning for Unsupervised Video Representation Learning

Anup Hiremath, Avidesh Zakhori; University of California Berkeley; Berkeley, California

Abstract

In this paper, we propose a multimodal unsupervised video learning algorithm designed to incorporate information from any number of modalities present in the data. We cooperatively train a network corresponding to each modality: at each stage of training, one of these networks is selected to be trained using the output of the other networks. To verify our algorithm, we train a model using RGB, optical flow, and audio. We then evaluate the effectiveness of our unsupervised learning model by performing action classification and nearest neighbor retrieval on a supervised dataset. We compare this triple modality model to contrastive learning models using one or two modalities, and find using all three modalities in tandem provides a 1.5% improvement in UCF101 classification accuracy, a 1.4% improvement in R@1 retrieval recall, a 3.5% improvement in R@5 retrieval recall, and a 2.4% improvement in R@10 retrieval recall as compared to using only RGB and optical flow, demonstrating the merit of utilizing as many modalities as possible in a cooperative learning model.

Introduction

The large amount of video data available in the modern internet has led to an increased emphasis on fast and effective ways to classify and learn from repositories of video data. However, training traditional supervised machine learning models on data requires the presence of labels, which are not easily available in most cases. This leads to the task of unsupervised representation learning: training a model that takes an unlabeled data sample as input and yields a vector that represents a sample's useful features as accurately and succinctly as possible. When a supervised task is provided, this pre-trained model can be used as the backbone to speed up training for the supervised task. The unsupervised model can be used as weight initialization for a supervised model, improving the accuracy of supervised learning models without the need for additional labelled data. Alternatively, the unsupervised model can work as an encoder, extracting the relevant features of video to train a fully connected layer on the supervised task.

This paper explores ways to leverage the different modalities available in video data to make a more effective unsupervised pre-training method. Similar to InfoNCE [1], our approach uses contrastive learning to create "pseudo-labels" to train the model with. In contrastive learning, a model is trained on unlabeled data to yield distinct outputs for samples, such as images or video clips, that are found to be dissimilar. In our case, we classify a pair of sample video clips as being similar or dissimilar using the other modalities in the sample. In one of our steps, we train an unsupervised model on RGB frame data by finding pairs of video clips that are similar in optical flow or audio space. This is effective because while two samples may differ in RGB feature space, they might prove similar in either the flow or audio space. InfoNCE

[1] uses only RGB data in a contrastive learning scheme and CoCLR [2] extends it to use RGB and optical flow. In this report, we extend CoCLR [2] to incorporate audio.




Figure 1. The overall training scheme. The green blocks represent networks that are being trained in a given stage, and the red blocks are the networks that are acting as the "oracle". These latter networks have their weights frozen, and are used to generate aggregated representations that can then be used to train the green network. After each modality network is trained relative to the others, the cycle repeats.

Related Work

Contrastive learning

The field of unsupervised learning is active, and the current state-of-the-art is contrastive learning, in which models are trained by learning to similarly represent augmentations of the same instance, or instances that are known to be similar, and vice versa [1][2][3][4][5][6]. InfoNCE [1] follows this formula, creating a positive pseudo-class by augmenting a sample multiple times, and a negative pseudo-class by augmenting other samples. The model is then trained to similarly represent pairs of samples both from the positive pseudo-class and distinctly represent other pairs. Momentum contrast, or MoCo [3] provides a sampling algorithm to run InfoNCE efficiently; negative examples are maintained in a queue in order to prevent having to read them from disk and augment them on every step. SimCLR [5] provides another sampling algorithm, in which each sample in a batch is augmented to produce N positive pairs, and then each positive pair is contrasted with the other $2(N - 1)$ augmented samples in the batch. DenseCL [4] uses momentum contrast, but rather than training based on a single global representation vector of each sample, it computes a representation vector of each feature of each sample, creating a dense representation to be contrasted. All of the above approaches focus on using contrastive learning for image data. In this paper, we build upon MoCo [3] for its speed and simplicity, both crucial when dealing with much larger video data.

Contrastive Video Learning

Much research has been done on adapting contrastive learning to the specific properties of video data [7][8][9][10][11].

CVRL [8] performs both spatial augmentations, consisting of standard image augmentations, and temporal augmentations, which involves selecting different timestamps of the same video. It then trains its model using the InfoNCE [1] loss. VideoMoCo [9] extends momentum contrast specifically for videos, by attenuating past samples in the queue to prefer those more recently added.

We train our model to yield similar representations invariant to temporal transformations; that is, we expect the same representation from two clips cut from the same video. Recent approaches such as Jenni et al. [10], instead train their model to specifically recognize temporal information. Behrmann et al. [11] use InfoNCE to train a model to recognize whether one clip takes place after another or vice versa, and use this specific temporal information to yield a representation learner. RSPNet [7] augments samples by playing them at different speeds, and trains a model to similarly represent samples from different videos that play at the same speed. None of these papers take into account other modalities represented in the data, and instead represent new ways of working with only RGB frame data.

Multimodal Video Learning

There are existing approaches that work with different modalities present in unlabeled video data [12][13][14][15][16][17]. One approach is to devise specific cross-modality tasks that can be used to construct a loss function. Arandjelović et al. [12] train a model to detect audio and video clips that correspond to the same time stamp. Korbar et al. [13] use this idea to train a representation learner to use in downstream supervised tasks. Piergiovanni et al. [14] construct many such cross-modality tasks, and incorporate the modalities into the formulation of the loss. Another approach, taken by XDC [15] and Asano et al. [16] is to use the different modalities to construct an effective clustering model and use the resulting clusters as pseudo-classes. GDT [17] treats the extraction of a modality from a video as an augmentation and incorporates augmentations into the loss function itself, thus learning to recognize samples from the same video even if the samples are of different modalities.

Our proposed approach is similar to that of CoCLR [2], which trains a neural network for each modality cooperatively, alternating the selected network to train using the other networks as an oracle. CoCLR [2] extracts RGB and dense optical flow, and uses a modified version of the InfoNCE [1] loss to train an RGB net and a flow net against one another. To train the RGB net, positive pairs for some given sample are extracted by getting the top k clips with the representation closest to that sample according to the optical flow network. CMA, or Cross-Modal Agreement, [6] uses RGB and audio instead, and rather than alternating one of the networks as the oracle, uses both the RGB and the audio representations of a sample to find similar samples. In this paper, we extend the alternating-oracle pattern of CoCLR [2] to three modalities: RGB, optical flow, and audio, thus increasing the information available to the models.

Approach

Our approach involves two stages: pre-training and training. In the pre-training step, we train three encoders using InfoNCE [1] with momentum contrast [3] for each one of our three modalities:

RGB, optical flow, and audio. In the training step, we use the pre-trained encoders in a cooperative contrastive learning scheme similar to CoCLR [2].

Pre-Training

Our algorithm, as illustrated in Figure 1, involves alternating selection of one of the modalities on which to train an encoder net, while freezing the encoders corresponding to other modalities. This is similar to the way CoCLR [2] alternates between RGB and optical flow, extended to three modalities. The frozen encoders act as an "oracle", determining which other samples in the dataset are most similar to a given sample. In order for those frozen encoders to serve as reasonable "oracles" to train the other encoders with, we need to first make sure they are trained to represent their respective modalities accurately. This is done using the InfoNCE [1] loss, which makes the encoder learn to recognize different random augmentations from the same video as being similar, leading it to encode only important features of the given clip. We use MoCo [3] to prevent having to fetch clips recompute augmentations for each time step.

Training

Once the encoders corresponding to each modality are trained we begin the cooperative training part of the algorithm, which is an extension of InfoNCE [1] and CoCLR [2] and the main contribution of this report. Our overall training scheme is shown in Figure 1: Given M pre-trained encoders corresponding to M modalities we select one encoder f for each stage of training, and use the other $M - 1$ encoders as oracles to train f . In this paper we use $M = 3$, with the modalities RGB, optical flow, and audio. With $M = 2$, this algorithm reduces to CoCLR [2], and with $M = 1$, this algorithm reduces further to InfoNCE [1]. Our proposed scheme outlined in Figure 1 has three stages corresponding to the three modalities; we summarize a single stage of training in Figure 2.

Similar to InfoNCE [1], our loss uses cross-entropy to train an encoder that encodes positive pairs similarly and all other pairs differently. The main difference between the single-modality pre-training algorithm and our multimodal algorithm is the way we select positive pairs. Rather than providing only a single positive pair (a_1, a_2) , we select k more positive matches for a_1 to form the positive set $\mathcal{P} = \{a_2, a_{p_1}, a_{p_2}, \dots, a_{p_k}\}$, as in CoCLR [2]. We use the other $M - 1$ encoders, which we refer to as *samplers*, to select these k positive samples. We freeze the weights of these samplers, colored red in Figure 2; they are only there to mine positive pairs for the encoder being trained, colored green in Figure 2. The outputs of these samplers is concatenated to form the *oracle* output.

Since running video clips through these sampler encoders to create positive and negative sets on each timestep is computationally expensive, we use momentum contrast [3]. This consists of two queues, in yellow in Figure 2, which are Queue B, consisting of past outputs of the encoder corresponding to the modality being trained, and Queue O, corresponding to past oracle outputs. The k most similar entries in Queue O corresponding to the current encoder output are matched to their corresponding entries in Queue B (colored dark green in Figure 2), and those entries are used to populate the positive set \mathcal{P} . The negative set \mathcal{N} consists of all of the remaining entries in P. Using this method greatly reduces the time spent in I/O, since no disk accesses are necessary

to compute the positive and negative sets for a given clip.

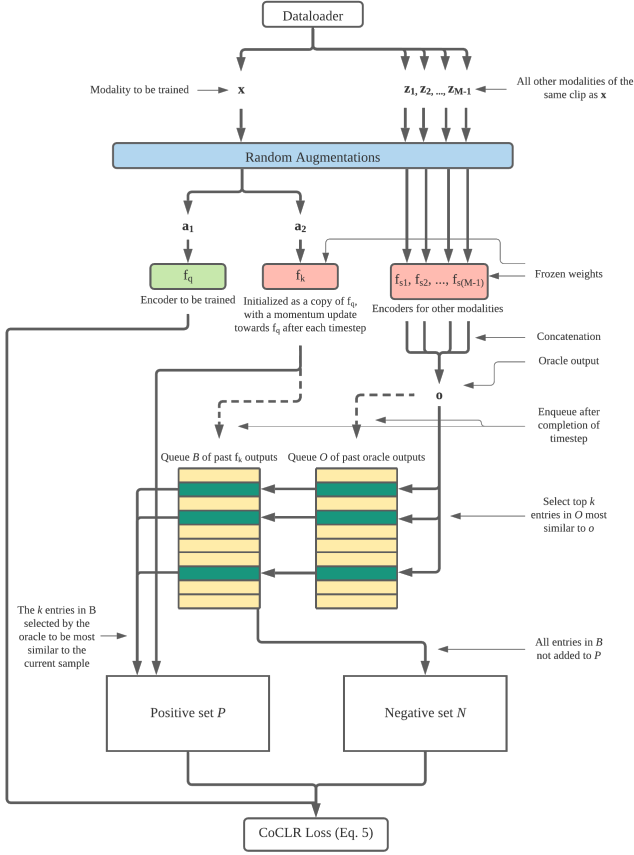


Figure 2. A single timestep of our entire multimodal contrastive learning algorithm for a single stage. We use Momentum Contrast [3] along with the CoCLR loss [2]. In our paper, $M = 3$, so there are two other modalities z_1 and z_2 , and two corresponding encoders f_{s1} and f_{s2} . The entire algorithm is computed in batches in order to speed up computation.

After the completion of one stage, a different modality is selected and the next stage begins. The recently trained encoder becomes one of the frozen samplers to make up the next oracle, and one of the frozen encoders from the previous oracle is selected to be trained next.

Experimental Setup

Modality Extraction

For unsupervised learning, we use the Kinetics-400 action classification dataset [18]. After downloading the data from YouTube, we extract the three modalities RGB, optical flow and audio. 300 RGB frames are extracted from each video and re-sized to 128 by 128. Optical flow is computed using Zach et al.’s TVL1 dense optical flow algorithm [19]. For each frame, flow in the x and y -directions is placed into the first and second channels of a 3-channel 128 by 128 image. Audio is processed into a mel-scaled spectrogram, saved as greyscale images. This allows us to use an out-of-the-box image classification backbone for our audio data.

Data Preparation

All of the data is packed into the Lightning Memory-Mapped Database (lmdb) format [20]: a serialized binary format designed to improve access speed for very large datasets. By using lmdb [20] we are able to greatly speed up the I/O portion of computation, allowing us to train our model on the full cleaned dataset of more than 200,000 video clips.

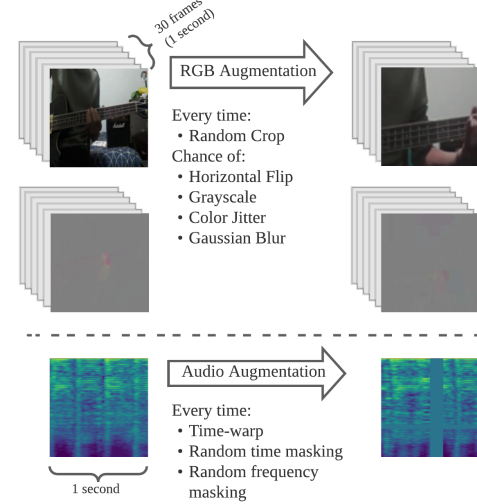


Figure 3. The augmentation types used on the three modalities in each sample: RGB frames, optical flow frames, and an audio spectrogram

Data Augmentations

In contrastive unsupervised representation learning algorithms such as InfoNCE [1] and SimCLR [5], the choice of data augmentation is crucial. Unlike in supervised tasks where data augmentation is an add-on to make models more robust, in contrastive learning, augmentations are an essential part of the learning process. In the absence of labels, the model is trained to recognize two inputs as different augmentations of the same clip. Without augmentations that accurately simulate the transformations and noise present in real data, the model will not be able to learn the essential features of its training data effectively.

We first perform temporal augmentation. At each iteration of the dataloader, we load the three pre-computed modalities of a video: RGB frames, optical flow, and audio. However, instead of simply sampling, augmenting, and returning a one-second clip from the video, it randomly samples two one-second clips, each from different parts of the same ten second video. The second of these clips is effectively a temporal augmentation of the first; it is a shift in time to a different part of the original video. Training the encoder to produce the same output invariant of this augmentation means different parts of the same video should yield the same representation.

After selecting two clips, spatial transformations are performed as shown in Figure 3. We use the same spatial augmentations for RGB and optical flow as CoCLR [2]. Every clip is randomly cropped, and then has a random chance of being horizontally flipped, color-jittered, Gaussian-blurred, and turned greyscale.

The spectrogram created from the audio portion of the two clips returned from temporal augmentation is randomly perturbed

in the time axis using a dense image warp. Then, time and frequency masking are applied, corresponding to a vertical and a horizontal band of the spectrogram being randomly replaced with the mean value of the spectrogram.

In the end, our dataset output at each step consists of six elements: three modalities from each of the two clips sampled from each video. The RGB and flow each are tensors of dimensions $batch_size \times 30 \times 3 \times 128 \times 128$, and the audio $batch_size \times 1 \times 128 \times 128$. This scheme helps maximize the amount of information provided to the model from each video clip.

Network Architectures

The model is set up in PyTorch, making use of the `torch.distributed` module to train data across multiple GPUs. For the RGB and Flow networks, we use the S3D network architecture [21] as the backbone. This was chosen because S3D is noted to be specifically well-suited to both RGB and Optical Flow inputs; likely because of this versatility, CoCLR [2] gets its best results using S3D. For the audio network, we use Resnet-18 [22], because it remains a simple, fast, and effective network architecture for image-based tasks. Ultimately, the choice of backbone is not the focus of this paper, since contrastive learning algorithms are agnostic of the backbone used. Both the 3D and 2D backbones can be easily swapped out as more effective network architectures emerge.

Training Implementation

We first pre-train three single-modality for 200 epochs unsupervised models using InfoNCE [1]: one on RGB data, one on audio data, and one on optical flow data. In order to evaluate the effectiveness of our multimodal training scheme, we then train a contrastive learning model using all three of these modalities. We train this model for two cycles of 3 stages each, with each stage being 200 epochs - every modality encoder was trained using the other encoders as the oracle for a total of 400 epochs. We set $k = 5$; that is, the oracle selects the top 5 most similar samples from its queue of past outputs.

This model was trained using an Azure compute instance with a Tesla V100 GPU, 110GB of RAM, an Intel Xeon E5-2690 CPU and a 1TB disk. The limiting factor in this was disk usage, as the data preparation process constantly ran into the 1TB limit. For our three-modality model, one training iteration took roughly 3 days, each 3-iteration stage took around 9 days, and the full two stages took around 3 weeks to train.

Results

We evaluate our unsupervised training method by testing our trained model on two supervised tasks. The first is action classification, in which our model is further trained on the UCF101 action classification dataset [23], and then evaluated against the UCF101 test set. The second is clip retrieval, in which the model retrieves the top N clips with the most similar encoder output to a given test clip; we evaluate whether any of the retrieved clips share the same class as the test clip. Though all three modalities are used in the training process, the effectiveness of our approach is evaluated using only the encoder dedicated to RGB. Through our cooperative learning method, our RGB encoder is able to infer information about audio and optical flow despite only being supplied RGB data during supervised evaluation.

Evaluation on Action Classification

We evaluate our model in the same way as past work on this topic [2] [6]: we fine-tune and evaluate our RGB encoder on the supervised action classification task UCF101 [23]. UCF101 consists of around 13,000 clips labelled in 101 different action classes, and is mostly a visual-only dataset; the lack of consistent audio present in most of its clips is the main reason we chose to perform our unsupervised training on Kinetics.

We train our supervised classification models in two ways:

1. Full training: no weights are frozen, and our encoder is trained end-to-end.
2. Linear probe: The entire encoder’s weights are frozen except for a final fully connected layer.

In both evaluations, we compare against approaches trained on Kinetics 400 for a more fair comparison with our model, since training using larger datasets will naturally result in improved classification accuracy. For instance, Elo [14], which is trained with YouTube8M, reaches a 93.8% accuracy, outperforming our model by 4.4%.

Table 1: UCF101 test accuracy after unsupervised training on Kinetics 400 and finetuning on UCF101

Model	Date	Modality			Top-1 Test Acc
		RGB	Flow	Audio	
InfoNCE [1]	2018	✓			79.5
CBT [24]	2019	✓			79.5
SpeedNet [25]	2020	✓			81.1
XDC [15]	2020	✓		✓	84.2
CMA [6]	2021	✓		✓	87.5
CoCLR [2]	2021	✓	✓		87.9
TCLR [26]	2022	✓	✓		88.2
STS [27]	2021	✓			89.0
GDT [28]	2021	✓		✓	89.3
RSPNet [7]	2021	✓			93.7
Ours	2022	✓	✓	✓	89.4

In the first evaluation, our unsupervised algorithm on Kinetics serves to initialize weights for our supervised training on UCF101. Since the entire model can be trained end-to-end using the supervised data, this provides the highest classification accuracy. Our model is compared against the state-of-the-art in Table 1, where we see that in full training, our model outperforms InfoNCE [1], which uses only RGB, by 9.9%, and CoCLR [2], which uses RGB and optical flow, by 1.5%. Our model also outperforms the models using only RGB and audio, XDC [15] and CMA [6], by 5.2% and 1.9% respectively, although neither of them use the alternating-oracle method used by our model and CoCLR.

Table 2: UCF101 test accuracy after unsupervised training on Kinetics 400 and linear probe on UCF101

Model	Date	Modality			Top-1 Test Acc
		RGB	Flow	Audio	
CBT [24]	2019	✓			54.0
MemDPC [29]	2020	✓			54.1
CoCLR [2]	2021	✓	✓		74.5
Ours	2022	✓	✓	✓	76.5

In Table 2, using the linear probe, our RGB encoder is truly acting as an encoder: our frozen RGB encoder encodes UCF101 clips as feature vectors and these feature vectors are used to train a single linear layer. This method of evaluation is less common, but provides better insight on the state of the unsupervised model before the introduction of labelled data. In Table 2, using the linear probe, our model outperforms CoCLR [2] by 2%. Thus, our improvement over CoCLR’s results shows that audio information can be incorporated successfully into a contrastive learning framework.

Evaluation on Retrieval

Quantitative Retrieval Results

We also evaluate unsupervised representation learners on the task of video retrieval. In this task, no supervised fine-tuning is done; the model being tested is trained purely using unsupervised data. Similar to classification, only the RGB encoder is used; the other encoders only serve to train the RGB encoder. First, for every sample in both the train and test splits of UCF101, we use the RGB encoder being evaluated to extract a feature vector representation of that sample. Then, for each feature vector from the test set, we look for the k vectors from the training set with the highest cosine similarity. We evaluate this retrieval by computing the recall $R@k$: the fraction of clips in the test set for which at least one of the k clips retrieved by our model is of the same action class as the test clip.

Table 3: UCF101 nearest-neighbors retrieval recall after unsupervised training on Kinetics-400

Model	Date	Modality			$R@1$	$R@5$	$R@10$
		RGB	Flow	Audio			
SpeedNet [25]	2020	✓			13.0	28.1	37.5
MemDPC [29]	2020	✓			20.2	40.4	52.4
STS [27]	2021	✓			38.3	59.9	68.9
CoCLR [2]	2021	✓	✓		44.5	60.6	68.4
Ours	2022	✓	✓	✓	45.9	64.1	70.8

In Table 3, we find that our model outperforms the state-of-the-art models. As expected, incorporating audio into our contrastive learning framework shows a 1.4% improvement in $R@1$ retrieval, a 3.5% improvement in $R@5$ retrieval, and a 2.4% improvement in $R@10$ retrieval performance over CoCLR[2].

Qualitative Retrieval Results

In Figure 4, we show two examples of how our RGB augmentations result in successful retrieval on UCF101 video clips. Though the clips in the both rows are oriented differently and are different colors, the random flips and color jitter applied as data augmentations enables the RGB encoder to recognize the videos as the same activity. In Figure 5, we show two examples of how incorporating optical flow affects retrieval results. In the top row, we can see that the model has learned to identify biking and horse riding as similar activities due to the similar motion in the videos, despite the differing backgrounds. Likewise, the clips in the bottom row all share a swinging motion despite the variety of activities. Both our model and CoCLR [2] yield similar results, since both models use RGB and optical flow in their approaches.

In Figure 6, we show two examples of how incorporating audio information improves retrieval results, which is the main con-

tribution of this report. In the top row, videos containing music are linked with one another, even if the scene and the motions within them are completely different. Music is easily identifiable on a spectrogram: notes and their harmonics show up as sharp, regularly spaced horizontal lines. In the second row, clips of drums and punches are retrieved together because of their similar percussive sounds. Since, unlike our model, CoCLR is never supplied audio data, it has not learned to associate clips containing music or percussive sounds with one another, and so it returns classes that are unrelated to the test clip.

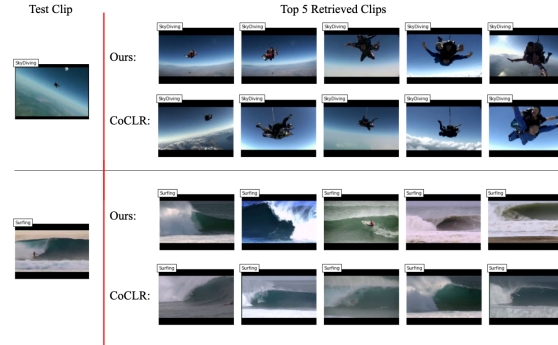


Figure 4. Qualitative effects of incorporating RGB information on UCF101 retrieval results. Our triple modality model is compared against CoCLR [2], both pre-trained on Kinetics-400.

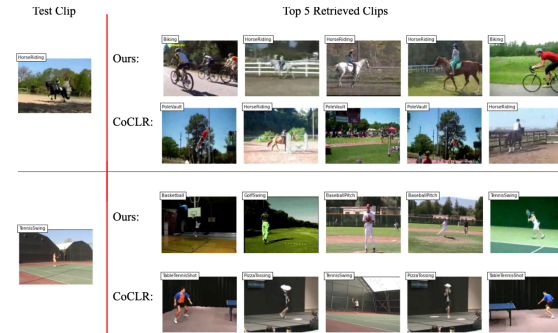


Figure 5. Qualitative effects of incorporating optical flow information on UCF101 retrieval results. Our triple modality model is compared against CoCLR [2], both pre-trained on Kinetics-400.



Figure 6. Qualitative effects of incorporating audio information on UCF101 retrieval results. Our triple modality model is compared against CoCLR [2], both pre-trained on Kinetics-400.

Conclusions and Future Work

Our approach extends the existing idea of cooperative contrastive video learning from unlabelled video data [2] to incorporate three modalities. We train a model using RGB, optical flow, and audio, and evaluate its performance against existing contrastive learning frameworks with fewer modalities to determine whether incorporating additional modalities using our scheme is effective. We train these models on Kinetics-400 and evaluate them on UCF101, using the tasks of action classification and nearest-neighbors retrieval. We found that in both tasks, our triple-modality model outperformed the contrastive learning models with fewer modalities and compared favorably to the state-of-the-art. Further improvements could be found in simply running our training scheme for more cycles, or by trying out other recent backbone architectures such as SlowFast [30]. Our training scheme can be easily extended to other modalities such as text: one could add a fourth encoder with an NLP backbone if subtitles or narration were available for video clips.

References

- [1] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding, 2018.
- [2] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning, 2021.
- [3] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning, 2019.
- [4] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training, 2021.
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations, 2020.
- [6] Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement, 2021.
- [7] Peihao Chen, Deng Huang, Dongliang He, Xiang Long, Runhao Zeng, Shilei Wen, Minghui Tan, and Chuang Gan. Rspnet: Relative speed perception for unsupervised video representation learning, 2021.
- [8] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning, 2021.
- [9] Tian Pan, Yibing Song, Tianyu Yang, Wenhao Jiang, and Wei Liu. Videomoco: Contrastive video representation learning with temporally adversarial examples, 2021.
- [10] Simon Jenni and Hailin Jin. Time-equivariant contrastive video representation learning, 2021.
- [11] Nadine Behrmann, Juergen Gall, and Mehdi Noroozi. Unsupervised video representation learning by bidirectional feature prediction, 2020.
- [12] Relja Arandjelović and Andrew Zisserman. Look, listen and learn, 2017.
- [13] Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative learning of audio and video models from self-supervised synchronization, 2018.
- [14] AJ Piergiovanni, Anelia Angelova, and Michael S. Ryoo. Evolving losses for unsupervised video representation learning, 2020.
- [15] Humam Alwassel, Dhruv Mahajan, Bruno Korbar, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering, 2020.
- [16] Yuki M. Asano, Mandela Patrick, Christian Rupprecht, and Andrea Vedaldi. Labelling unlabelled videos from scratch with multi-modal self-supervision, 2021.
- [17] Mandela Patrick, Yuki M. Asano, Polina Kuznetsova, Ruth Fong, João F. Henriques, Geoffrey Zweig, and Andrea Vedaldi. On compositions of transformations in contrastive self-supervised learning, 2021.
- [18] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset, 2017.
- [19] Christopher Zach, Thomas Pock, and Horst Bischof. A duality based approach for realtime tv-l1 optical flow. volume 4713, pages 214–223, 09 2007.
- [20] Howard Chu. Lightning memory-mapped database manager (lmdb), 2011.
- [21] Da Zhang, Xiyang Dai, Xin Wang, and Yuan-Fang Wang. S3d: Single shot multi-span detector via fully 3d convolutional networks, 2018.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.
- [23] Khuram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild, 2012.
- [24] Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Learning video representations using contrastive bidirectional transformer, 2019.
- [25] Sagie Benaim, Ariel Ephrat, Oran Lang, Inbar Mosseri, William T. Freeman, Michael Rubinstein, Michal Irani, and Tali Dekel. Speednet: Learning the speediness in videos, 2020.
- [26] Ishan R. Dave, Rohit Gupta, Mamshad Nayeem Rizve, and Mubarak Shah. TCLR: temporal contrastive learning for video representation. *CoRR*, abs/2101.07974, 2021.
- [27] Jiangliu Wang, Jianbo Jiao, Linchao Bao, Shengfeng He, Wei Liu, and Yunhui Liu. Self-supervised video representation learning by uncovering spatio-temporal statistics. *CoRR*, abs/2008.13426, 2020.
- [28] Mandela Patrick, Yuki Asano, Polina Kuznetsova, Ruth Fong, João F. Henriques, Geoffrey Zweig, and Andrea Vedaldi. Multi-modal self-supervision from generalized data transformations, 2021.
- [29] Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning, 2020.
- [30] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition, 2019.

Author Biography

Anup Hiremath received his BS (2020) and his MS (2022) in Electrical Engineering and Computer Science from UC Berkeley. Since then, he works as a software engineer at Amazon in San Francisco, California. His research is in the Video and Image Processing Lab at UC Berkeley, and focuses on deep learning and computer vision.