

CoCon: Cooperative-Contrastive Learning

Nishant Rai¹, Ehsan Adeli¹, Kuan-Hui Lee², Adrien Gaidon², Juan Carlos Nieves¹
¹Stanford University ²Toyota Research Institute

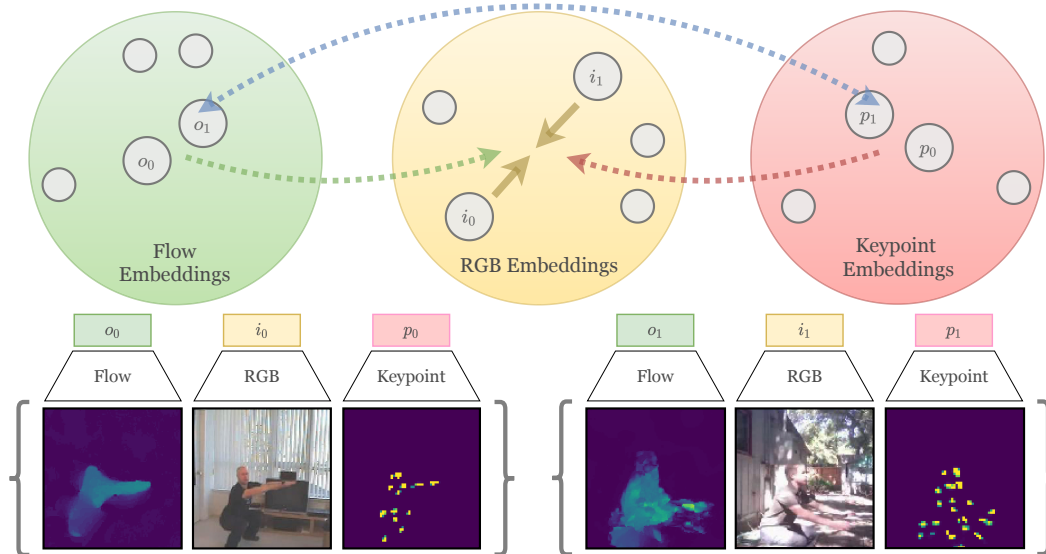


Figure 1: Given a pair of instances (e.g. people doing squats) and corresponding multiple views, features are computed using view-specific deep encoders f 's. Different instances may have contrasting similarities in different views. For instance, V_0 (left) and V_1 (right) have similar optical-flow $o = f_{flow}$ and pose keypoints (keypoint) $p = f_{keypoint}$ features but their image $i = f_{rgb}$ features are far apart. CoCon leverages these inconsistencies by encouraging the distances in all views to become similar. High similarity of o_0, o_1 and p_0, p_1 nudges i_0, i_1 towards each other in the RGB space.

Abstract

Labeling videos at scale is impractical. Consequently, self-supervised visual representation learning is key for efficient video analysis. Recent success in learning image representations suggest contrastive learning is a promising framework to tackle this challenge. However, when applied to real-world videos, contrastive learning may unknowingly lead to separation of instances that contain semantically similar events. In our work, we introduce a cooperative variant of contrastive learning to utilize complementary information across views and address this issue. We use data-driven sampling to leverage implicit relationships between multiple input video views, whether observed (e.g. RGB) or inferred (e.g. flow, segmentation masks, poses). We are one of the firsts to explore exploiting inter-instance relationships to drive learning. We experimentally evaluate our representations on the downstream task of action recognition.

Our method achieves competitive performance on standard benchmarks (UCF101, HMDB51, Kinetics400). Furthermore, qualitative experiments illustrate that our models can capture higher-order class relationships. The code is available at <http://github.com/nishantrai18/CoCon>.

1. Introduction

There has recently been a surge in interest for approaches utilizing self-supervised methods for visual representation learning. Recent advances in visual representation learning have demonstrated impressive performance compared to their supervised counterparts [3, 14]. Fresh development in the video domain have attempted to make similar improvements [10, 16, 25, 35].

Videos are a rich source for self-supervision, due to the inherent temporal consistency in neighboring frames. A natural approach to exploit this temporal structure is pre-

dicting future context as done in [10, 16, 25, 27]. Such approaches perform future prediction in mainly two ways: (1) predicting a reconstruction of future frames [25, 27, 39], (2) predicting features representing the future frames [10, 16]. If the goal is learning high-level semantic features for other downstream tasks, then complete reconstruction of frames is unnecessary. Inspired by developments in language modelling [29], recent work [41] propose losses that only focus on the latent embedding using frame-level context. One of the more recent approaches [10] propose utilizing spatio-temporal context to learn meaningful representations. Even though such developments have led to improved performance, the quality of the learned features is still lagging behind that of their supervised counterparts.

Due to the lack of labels in self-supervised settings, it is impossible to make direct associations between different training instances. Instead, prior work has learned associations based on structure, either in the form of temporal [10, 20, 23, 26, 44] or spatial proximity [10, 18, 20, 30] of patches extracted from training images or videos. However, the contrastive losses utilized enforce similarity constraints between instances from same videos while pushing instances from other videos far away even if they represent the same semantic content. This inherent drawback forces learning of features with limited semantic knowledge and encourage performing low-level discrimination between different videos. Recent approaches suffer from this restriction leading to poor representations.

The idea of utilizing multiple views of information has been a well-established one with roots in human perception [4, 15]. It’s argued that useful higher order semantics are present throughout different views and are consistent across them. At the same time, different views provide complementary information which can be utilized to aid learning in other views. Multi-view learning has been a popular direction [35, 40] utilizing these traits to improve representation quality. Recent approaches learn features utilizing multiple views with the motivation that information shared across views has valuable semantic meaning. A majority of these approaches directly utilize core ideas such as contrastive learning [31] and mutual information maximization [2, 24, 46]. Although the fusion of views leads to improved representations, such approaches also utilize contrastive losses, consequently suffering from the same drawback of low-level discrimination between similar instances.

We propose Cooperative Contrastive Learning (CoCon), which overcomes this shortcoming and leads to improved visual representations. Our main motivation is that each view sees a specific pattern, which can be useful to guide other views and improve representations. Our approach utilizes inter-view information to avoid the drawback of discriminating similar instances discussed earlier. To this end, each view sees a different aspect of the videos, allowing it to

suggest potentially similar instances to other views. This allows us to infer implicit relationships between instances in a self-supervised multi-view setting, something which we are the first to explore. These associations are then used in order to learn better representations for downstream applications such as video classification and action recognition. Fig. 1 shows an overview of CoCon. It is worth noting that although CoCon utilizes building blocks currently used in self-supervised representation learning, it is applicable to other tasks utilizing contrastive learning and be used in conjunction with other recently proposed methods.

We use ‘freely’ available views of the input such as RGB frames and Optical Flow. We also explore the benefit of using high-level inferred semantics as additional noisy views, such as human pose keypoints and segmentation masks generated using off-the-shelf models [45]. These views are not independent, as they can be derived from the original input images. However, they are complementary and lead to significant gains, demonstrating CoCon’s effectiveness even with noisy related views. The extensible nature of our framework and the ‘freely’ available views used make it possible to use CoCon with any publicly available video dataset and other contrastive learning approaches.

2. Related Work

Self-supervised Learning from images Recent approaches have tackled image representation learning by exploiting color information [22, 47] and spatial relationships [30, 34], where relative positions between image patches are exploited as supervisory signals. Several approaches apply self-supervision to super-resolution [6, 19] or even to multi-task [5] and cross-domain [33] learning frameworks.

Self-supervised Learning from videos Multiple approaches [10, 16, 25, 27, 39] perform self-supervision through ‘predicting’ future frames. However, the term ‘predicting’ is overloaded, as they do not directly predict and reconstruct frames but instead operate on latent representations. This ignores stochasticity of frame appearance, e.g., illumination changes, camera motion, appearance changes due to reflections and so on, allowing the model to focus on higher-order semantic features. Recent work [10, 40] utilize Noise Contrastive Estimation to perform prediction of the latent representations rather than the exact future frames, vastly improving performance. Yet, another class of proxy tasks are based on temporal ordering of frames [28, 44]. Temporal coherence [17, 43] and 3D puzzle [20] were used as proxy loss to exploit spatio/temporal structures.

Multi-view learning Multiple views of videos are rich sources of information for self-supervised learning [35, 40, 42]. Two stream networks for action recognition [37] have led to many competitive approaches, which demonstrate using even derivable views such as optical flow helps improve performance considerably. There have been approaches

[26, 35, 40, 42] utilizing diverse views, sometimes derivable from one other, to learn better representations. However, these approaches utilize inter-view links by maximizing mutual information between them. Although this leads to improved performance, we believe the rich inter-view linkages can be utilized more effectively by utilizing them to uncover implicit relationships between instances.

Multi-View Self-supervised learning Multiple recent approaches [1, 11, 12, 32] have tackled the challenge of multi-modal self-supervised learning achieving impressive performance. However, these approaches suffer from the same drawback of discriminating between similar instances, leaving potential to benefit from inter-sample relationships.

Most approaches above perform self-supervision using positive and negative pairs mined through structural constraints, e.g., temporal and spatial proximity. Although this results in representations that capture some degree of semantic information, it incorrectly leads to treating similar actions differently due to the inherent nature of their pair-mining. For instance, clip pairs in different videos are considered negatives, even if they represent the same action. We argue that utilizing different views and inter-instance relationships to propose positive pairs to aid training can lead to improvement of all views simultaneously.

3. Method

We describe cooperative contrastive learning (CoCon) and intuition behind our designs in this section. Additional details regarding architecture and implementation are present in the appendix. In the following sections, we build our framework borrowing the learning framework present in [10] which learns video representations through spatio-temporal contrastive losses. It should be noted that even though we use this particular self-supervised backbone in our experiments, our approach is not restricted by the choice of the underlying self-supervised task. CoCon can be used in conjunction with any other frameworks currently present and allow them to be extended to a multi-view setting.

A video V is a sequence of T frames (not necessarily RGB images) with resolution $H \times W$ and C channels, $\{\mathbf{i}_1, \mathbf{i}_2, \dots, \mathbf{i}_T\}$, where $\mathbf{i}_t \in \mathbb{R}^{H \times W \times C}$. Assume $T = N * K$, where N is the number of blocks and K denotes the number of frames per block. We partition a video clip V into N disjoint blocks $V = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$, where $\mathbf{x}_j \in \mathbb{R}^{K \times H \times W \times C}$ and a non-linear encoder $f(\cdot)$ transforms each input block x_j into its latent representation $z_j = f(x_j)$. An aggregation function, $g(\cdot)$ takes a sequence $\{z_1, z_2, \dots, z_j\}$ as input and generates a context representation $c_j = g(z_1, z_2, \dots, z_j)$. In our setup, $z_j \in \mathbb{R}^{H' \times W' \times D}$ and $c_j \in \mathbb{R}^D$. D represents the embedding size and H' , W' represent down-sampled resolutions as different regions in z_j represent features for different spatial locations. We define $z'_j = Pool(z_j)$ where $z'_j \in \mathbb{R}^D$ and $c = F(V)$ where

$$F(\cdot) = g(f(\cdot)).$$

Similar to [10], we create a prediction task involving predicting z of future blocks. Details are provided in the appendix. For multiple views, we define $c_v = F_v(V_v)$, where V_v, c_v and F_v represent the input, context feature and composite encoder for view v respectively.

Contrastive Loss Noise Contrastive Estimation (NCE) [9, 29, 31] constructs a binary classification task where a classifier is fed with real and noisy samples with the training objective being distinguishing them. Similar to [10, 31], we use an NCE loss over our feature embeddings described in Eq 1. $z_{i,k}$ represents the feature embedding for the i^{th} time-step and the k^{th} spatial location. Recall $z_j \in \mathbb{R}^{H' \times W' \times D}$ which preserves the spatial layout. We normalize $z_{i,k}$ to lie on the unit hypersphere. Eq 1 is a cross-entropy loss distinguishing one positive pair from all the negative pairs present in a video. We use temperature $\tau = 0.005$ in our experiments. In a batch setting with multiple video clips, it is possible to have more inter-clip negative pairs.

To extend this to multiple views, we utilize different encoders ϕ_v for each view v . We train these encoders by utilizing \mathcal{L}_{cpc} for each of them independently, giving us, $\mathcal{L}_{cpc} = \sum_v \mathcal{L}_{cpc}^v$

$$\mathcal{L}_{cpc} = - \sum_{i,k} \left(\log \frac{\exp(\tilde{z}_{i,k} \cdot z_{i,k} / \tau)}{\sum_{j,m} \exp(\tilde{z}_{i,k} \cdot z_{j,m} / \tau)} \right) \quad (1)$$

Cooperative Multi-View Learning Recent approaches [12, 35, 40] tackle multi-view self-supervised learning by maximizing mutual information across views. They involve using positive and negative pairs generated using structural constraints, e.g., spatio-temporal proximity in videos [10, 11, 35, 40]. Although such representations capture semantic content, they unintentionally encourage discriminating video clips containing semantically similar content due to the inherent nature of pair generation, i.e. video clips from different videos are negatives. We utilize inter-instance relationships to alleviate some of these issues.

We soften this constraint by indirectly deriving pair proposals using different views. Such a co-operative scheme benefits all models as each individual view gradually improves. Better models are able to generate better proposals, improving performance of all views creating a positive feedback loop. Our belief is that significant semantic features should be universal across views, therefore, potential incorrect proposals from one view should cancel out through proposals from other views.

We achieve the above by computing view-specific distances and synchronizing them across all views. We enforce a consistency loss between distances from each view. Looking at it from another perspective, we are encouraging relationships between instances to be the same across views i.e. similar pairs in one view should be a similar pair in




Figure 2: Examples for each view. From top to bottom - RGB, Flow, SegMasks and Poses. Note the prevalence of noise in a few samples, specially SegMasks; There are multiple other instances where Poses, SegMasks are noisy but have not been shown here.

other views as well. Treating this as inter-view graph regularization, we create a graph similarity matrix W_v of size $K \times K$, using some distance metric. We represent our distance metric by $\mathcal{D}(\cdot)$. In our experiments, we use the cosine distance which translates to $W_{ab}^v = z_z \cdot z_b$.

Assume h_v^a denotes the representation for the v^{th} view of instance a . In our experiments, we use $h = z'$ giving us block level features. Our resultant loss becomes the inconsistency between similarity matrices across views. The resultant graph regularization loss becomes $\sum_{v_0, v_1} \|W^{v_0} - W^{v_1}\|$ which is simplified in Eq 2.

Building on top of our earlier intuition, in order to have sensible proposals, we need to have discriminative scores, i.e. we should have both positive ($\mathcal{D} \rightarrow 0$) and negative ($\mathcal{D} \rightarrow 1$) pairs. To promote well distributed distances, we utilize the hinge loss described in Eq 3.

\mathcal{L}_{sim} is the hinge loss, where the first term pushes representations of the same instance in different views closer; while the second term pushes different instances apart. Since the number of structural negative pairs are much larger than the positives, we introduce μ in order to balance the loss weights. We choose μ such that the first and second components contribute equally to the loss.

$$\mathcal{L}_{sync} = \sum_{v_0, v_1} \sum_{a, b} \left(\mathcal{D}(h_{v_0}^a, h_{v_0}^b) - \mathcal{D}(h_{v_1}^a, h_{v_1}^b) \right)^2 \quad (2)$$

View	Random	\mathcal{L}_{cpc}	\mathcal{L}_{sim}^{cpc}	\mathcal{L}_{sync}^{cpc}	\mathcal{L}_{cocon}
RGB	46.7	63.7	66.0	62.7	67.8
Flow	65.3	69.8	71.4	69.2	72.5

Table 1: Impact of losses on performance of models when jointly trained with RGB and Flow. CoCon i.e. \mathcal{L}_{total} (67.8) comfortably improves performance over CPC i.e. \mathcal{L}_{cpc} (63.7). $\mathcal{L}_y^x = \mathcal{L}_x + \lambda \mathcal{L}_y$ where $\lambda = 10.0$ for this experiment

Method	Pretrain	RGB		Flow ²	
		UCF	HMDB	UCF	HMDB
Random		46.7	20.6	65.3	31.2
CPC	K400	68.6	35.5	69.8	40.8
CoCon	UCF	67.8	37.7	72.5	44.1
CoCon	K400	72.1	46.5	71.8	44.2

Table 2: Impact of pre-training comparison. CoCon demonstrates a consistent improvement in both RGB and Flow.

Method	RGB		Flow		PoseHM		SegMask	
	UCF	HMDB	UCF	HMDB	UCF	HMDB	UCF	HMDB
Random	46.7	20.6	65.3	31.2	51.7	33.0	42.7	26.3
CPC	63.7	33.1	71.2	44.6	56.4	42.0	53.7	32.8
CoCon	71.0	39.0	74.5	45.4	58.7	42.6	55.8	34.0

Table 3: Impact of co-training on views. CoCon is jointly trained with four modalities (RGB, Flow, PoseHM, & SegMask).

$$\mathcal{L}_{sim} = \sum_{v_0, v_1} \sum_a \mathcal{D}(h_{v_0}^a, h_{v_1}^a) + \mu \sum_{a \neq b} \max(0, 1 - \mathcal{D}(h_{v_0}^a, h_{v_1}^b)) \quad (3)$$

Note that \mathcal{L}_{sim} entangles different views together. An alternative would be defining such a loss individually for each view. However, diversity is inherently encouraged through \mathcal{L}_{cpc} , and interactions between views have the side-effect of increasing their mutual information (MI), which leads to improved performance [35, 40].

We combine the above losses to get our cooperative loss, $\mathcal{L}_{coop} = \mathcal{L}_{sync} + \alpha \cdot \mathcal{L}_{sim}$. We use $\alpha = 1.0$ for our experiments and observe roughly similar performance for different values of α . The overall loss of our model is given by $\mathcal{L}_{cocon} = \mathcal{L}_{cpc} + \lambda \cdot \mathcal{L}_{coop}$. \mathcal{L}_{cpc} encourages our model to learn good features for each view, while \mathcal{L}_{coop} nudges it to learn higher-level features using all views while respecting the similarity structure across them.

4. Experiments

The goal of our framework is to learn video representations which can be leveraged for video analysis tasks. Therefore, we perform experiments validating the quality of our representations. We measure downstream action classification to objectively measure model effectiveness and ana-

Action Class	CoCon	CPC
PlayCello	PlaySitar, PlayTabla, PlayDhol Surfing, Skijet	N/A
Skiing		Surfing
HammerThrow	BaseballPitch, ThrowDiscus, Shotput	N/A
BrushTeeth	ApplyLipstick, EyeMakeup, ShaveBeard	ApplyLipstick

Table 4: Nearest consistent semantic classes. Individually trained views (*CPC*) do not have consistent neighbors across views, leading to empty results (N/A) for ‘PlayingCello’ and ‘HammerThrow’. While views trained using CoCon show consistency across views, leading to sensible relationships e.g. ‘HammerThrow’ related to other classes involving throwing.

lyze impact of our designs through controlled ablation studies. We also conduct qualitative experiments to gain deeper insights into our approach. In this section, we briefly go over our experiment framework. Additional details and discussions for each component are provided in the appendix.

Datasets Our approach is a self-supervised learning framework for any dataset with multiple views. However, we discuss its relevance to video action classification in our experiments. We focus on human action datasets i.e. UCF101, HMDB51 and Kinetics400. UCF101 contains 13K videos spanning over 101 human action classes. HMDB51 contains 7K video clips mostly from movies for 51 classes. Kinetics-400 (K400) is a large video dataset with 306K video clips from 400 classes.

Views We utilize different views in our experiments. For Kinetics-400, we learn encoders for RGB and Optical Flow. We use Farneback flow (FF) [7] instead of the commonly used TVL1-Flow as it is quicker to compute lowering our computation budget. Although FF leads to lower performance compared to TVL1, the essence of our claims remain unaffected. For UCF101 and HMDB51, we learn encoders for RGB, TVL1 Optical Flow, Pose Heatmaps (PoseHMs) and Human Segmentation Masks (SegMasks). A few visual samples for each view are provided in 2. PoseHMs and SegMasks are generated using an off-the-shelf detector [45] without any form of pre/post-processing.

Implementation Details We choose a 3D-ResNet similar to [10, 13] as the encoder $f(\cdot)$. We choose $N = 8$ and $K = 5$ in our experiments. We subsample the input by uniformly choosing one out of every 3 frames. Our predictive task involves predicting the last three blocks using the first five blocks. We use standard data augmentations during training whose details are provided in the appendix. We train our models using Adam [21] optimizer with an initial learning rate of 10^{-3} , decreased upon loss plateauing. We use 4 GPUs with a batch size of 16 samples per GPU. Multiple spatio-temporal samples ensure sufficient negative examples despite the small batch size used for training.

Action Classification We measure the effectiveness of our learned representations using the downstream task of action classification. We follow the standard evaluation protocol of using self-supervised model weights as initializa-

# Views	RGB		Flow	
	UCF	HMDB	UCF	HMDB
2	67.8	37.7	72.5	44.1
4	71.0	39.0	74.5	45.4

Table 5: Impact of performance on varying views. A consistent improvement can be seen with more views despite the prevalent noise in PoseHM and SegMasks.

tion for supervised learning. The architecture is then fine-tuned end-to-end using class label supervision. We finally report the fine-tuned accuracies on UCF101 and HMDB51. While fine-tuning, we use the learned composite function $F(\cdot)$ in order to generate context representations for the video blocks. The context feature is further passed through a spatial pooling layer followed by a fully-connected layer and a multi-way softmax for action classification.

4.1. Quantitative Results

We analyze various aspects of CoCon through ablation studies, experiments on multiple datasets, controlled variation of views and comparison to comparable methods. We objectively evaluate model performance using downstream classification accuracy as a proxy for learned representation quality. Pre-training is performed on either UCF101 or Kinetics400. We propose two baselines for comparison. (1) *Random* - random initialization of weights (2) *CPC* - self-supervised training utilizing only \mathcal{L}_{cpc} ; which is effectively individual training of views. *CPC* serves as a critical baseline to measure the benefits of multi-view training as opposed to individual training.

Ablation Study We have motivated the utility of our various loss components. We now perform experiments to quantify the impact of each. The pre-training dataset used is the 1st split of UCF101, and downstream classification accuracy is computed on the same. Table 1 summarizes the results of our experiment. As expected, all cross-view approaches comfortably perform better than *CPC*; demonstrating the utility of multi-view training.

Using \mathcal{L}_{sync}^{cpc} leads to no performance improvements, as only using \mathcal{L}_{sync} leads to the model collapsing by squashing all \mathcal{D} scores to have similar values, thus necessitating \mathcal{L}_{sim} to counter-balance this tendency. \mathcal{L}_{sim}^{cpc} leads to improved performance wrt \mathcal{L}_{cpc} as it learns better features by effectively maximizing mutual information between views. CoCon i.e. $\mathcal{L}_{cococon}$ achieves the same by also regularizing manifolds across views, leading to even better performance across all views. The important comparison to observe is between \mathcal{L}_{sim}^{cpc} and $\mathcal{L}_{cococon}$. As \mathcal{L}_{sim}^{cpc} is the most similar baseline to other multi-view approaches, e.g., CMC [40]. However, we argue this baseline is even stronger as it in-

volves both single-view and multi-view components compared to [40], which only uses a contrastive multi-view loss to learn representations.

Effect of Datasets A critical benefit of self-supervised approaches is the ability to run on large unlabelled datasets. To simulate such a setting, we perform pre-training using UCF101 or Kinetics400¹ without labels utilizing the 1st splits of UCF101 and HMDB51 for evaluation. Table 2 confirms pre-training with a larger dataset leads to better performance. It is also worth noting that CoCon pre-trained with UCF101 outperforms *CPC* trained on Kinetics400 even though CoCon on UCF101 uses only around 10% data compared to Kinetics. Further demonstrating the potential of utilizing multiple views as opposed to training with larger and diverse datasets.

When comparing the *Random* baseline and CoCon pre-trained on Kinetics400, we observe higher performance gains for RGB (+25.4%) compared to Optical-Flow (+6.9%). We argue this is due to higher variance and complexity of RGB compared to Flow, allowing a randomly initialized network to perform relatively better with Flow. While comparing our approach with *CPC*, we again observe higher gains in RGB (+4.1%) compared to Flow (+2.7%). This can be explained by the potential capability of RGB to capture flow-like features when learned jointly.

Effect of cooperative training We compare benefits of cooperative training with varying views. We look at co-training of RGB, Flow, SegMasks and PoseHMs. Recall that these additional views are generated using off-the-shelf models without any additional post-processing. Even though they are somewhat redundant i.e. Flow, PoseHM, SegMasks are actually derived from RGB Images; using them simultaneously still leads to a large performance increase. We also note that although SegMasks and PoseHMs are sparse low-dimensional features, they still help improve performance across all views.

Table 3 summarizes downstream action recognition performance of each view under different approaches. We see improved performance with increase in the number of views used. Consistent gains for views such as Flow, SegMasks, PoseHM, which are not as expressive as RGB points towards extraction of higher-order features even from low dimensional inputs. We observe PoseHM and SegMask have lower performance gains when evaluated on HMDB51. This can be attributed to the large degree of noise in PoseHMs and SegMasks for HMDB51. HMDB is a challenging and diverse dataset, leading to poor predictions from our off-the-shelf detector. In conclusion, the benefits of joint training are apparent as CoCon leads to a performance improvement for all the views involved.

¹ Optical Flow used for Kinetics400 is Farneback Flow; as opposed to TVL1 Flow for UCF101 and HMDB51. This difference in pre-training and fine-tuning modalities leads to less than expected performance gains.

Effect of additional views CoCon hinges on the assumption that multi-view information helps in improving overall representation quality. To verify our hypothesis, we study co-training with different number of views. We consider two scenarios, 1) Joint training of RGB and Flow streams, and 2) Joint training of RGB, Flow, SegMasks and PoseHMs. Table 5 shows a consistent increase across views when increasing the number of views used during training. We should note that both SegMasks and PoseHMs contain significant noise as the off-the-shelf models incorrectly detects and misses humans in numerous videos. However, we see a consistent mutual increase in performance for all the involved views despite the prevalence of noise.

Comparison with comparable approaches We summarize comparisons of CoCon with comparable state-of-the-art approaches in Table 6. CoCon-Ensemble refers to an ensemble of models for all the involved views. We observe a few major trends, (1) When pre-training on UCF101, using multiple views allows us to outperform the nearest comparable approach by around 10.4%. This demonstrates the potential of cooperatively utilizing multiple views to learn representations. (2) We see considerable gains while training on Kinetics400 as well, however, the increase is smaller compared to UCF101. We argue the reasons are, a) we only utilize two views for co-training. b) the flow we utilize for Kinetics400 is Farneback Flow instead of TVL1 flow used for UCF101 and HMDB51. (3) Our method comfortably outperforms recent multi-view approaches consistently on UCF101 and HMDB51. (4) An interesting observation is that using multiple views of a small dataset (UCF101) performs better (71.0%) than pre-training on a large dataset, Kinetics400 (68.2%). This suggests that utilizing different views can be better than merely training on larger datasets.

Comparison with recent approaches A few very recent approaches [1, 11, 12, 32] have tackled multi-modal self-supervised achieving impressive performance. CoCon differs from them as it considers inter-instance relationships to aid learning in addition to relationships between views. Due to resource constraints, it was not possible to have a fair comparison due to the significant difference in the amount of GPUs, number of epochs trained and the backbones used. However, we hope our carefully constructed experiments given earlier provide deeper insights into CoCon’s benefits even with lower resource requirements.

4.2. Qualitative Results

We motivate CoCon arguing about the benefits of preserving similarities across view-specific spaces. We observe respecting structure across views results in emergence of higher-order semantics without additional supervision e.g. sensible class relationships and good feature representations. Jointly training with views known to perform well for video action understanding allows us to learn good video

Method	Resolution	Backbone	# Views	Pre-train	UCF101	HMDB51
Random Initialization	128 × 128	ResNet18	1		46.7	20.6
ImageNet [36]	224 × 224	VGG-M-2048	1	ImageNet	73.0	40.5
Shuffle and Learn [28]	227 × 227	CaffeNet	1	UCF-HMDB	50.2	18.1
OPN [23]	80 × 80	VGG-M-2048	1	UCF-HMDB	59.8	23.8
DPC [10]	128 × 128	ResNet18	1	UCF101	60.6	-
VGAN [42]	N/A	C3D	2	Flickr [42]	52.1	-
LT-Motion [26]	N/A	RNN [26]	2	NTU	53.0	-
Cross and Learn [35]	224 × 224	CaffeNet	2	UCF101	58.7	27.2
Geometry [8]	N/A	CaffeNet	2	UCF101	55.1	23.3
CMC [40]	128 × 128	CaffeNet	3	UCF101	59.7	26.1
CoCon - RGB	128 × 128	ResNet18	4	UCF101	70.5	38.4
CoCon - Ensemble	128 × 128	ResNet18	4	UCF101	82.4	52.0
3D-RotNet [18]	112 × 112	ResNet18	1	Kinetics	62.9	33.7
DPC [10]	128 × 128	ResNet18	1	Kinetics	68.2	34.5
CoCon - RGB	128 × 128	ResNet18	2	Kinetics	71.6	46.0
CoCon - Ensemble	128 × 128	ResNet18	2	Kinetics	78.1	52.0
ST-Puzzle [20]	224 × 224	ResNet18	1	Kinetics	65.8	33.7
DPC [10]	224 × 224	ResNet34	1	Kinetics	75.7	35.7
CoCon - RGB	224 × 224	ResNet34	2	Kinetics	79.1	48.5
CoCon - Ensemble	224 × 224	ResNet34	2	Kinetics	82.0	53.1

Table 6: Comparison of classification accuracies on UCF101 and HMDB51, averaged over all splits.

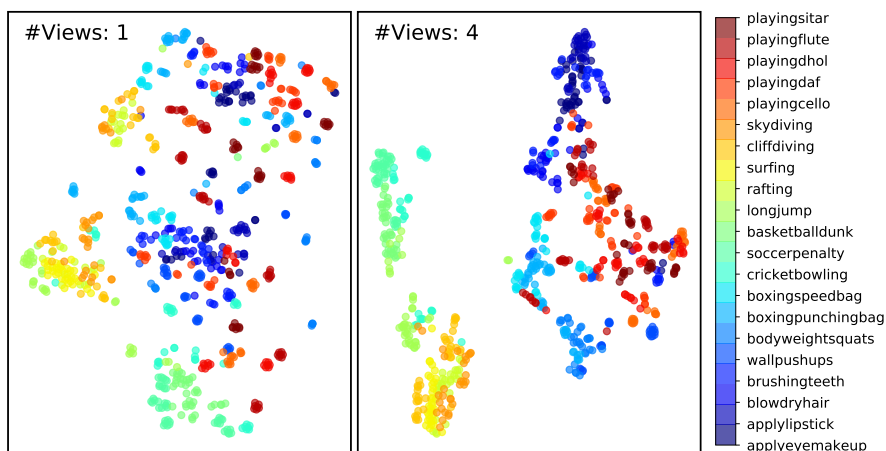


Figure 3: t-SNE visualization of RGB features from *CPC* (left) and *CoCon* (right) trained with 4 views. The color mapping for each category represents the relationships between action classes, e.g., Red: Instruments; Yellow: Water Sports; Light-blue: Physical Acts; Blue: Makeup-Hygiene. More meaningful clusters are formed using *CoCon*; signifying the ability of *CoCon* to align different yet semantically-related classes without any additional supervision.




Figure 4: Differences between class-wise accuracy for *CoCon* vs *CPC*. Only extreme classes are displayed. Blue - Gains; Red - Loss

representations, consequently, imparting unexpected side effects such as action alignment across videos. We discuss various experiments and results to support these claims.

t-SNE Visualization We explore t-SNE visualizations of our learned representations on the 1st test split of UCF101 extracted using $F(\cdot)$. For clarity, only 21 action classes are displayed. We loosely order the action classes according to their relationships. Classes having similar colors are seman-

tically similar. We can roughly observe the following broad categories present in the mentioned classes: Playing Instruments, Water Sports, Physical Sports, Physical Activities, Makeup-Hygiene. Results are displayed in Fig 3. Although we operate in a self-supervised setting, *CoCon* is able to uncover deeper semantic features allowing us to uncover inter-class relationships. We can see a much more concise and consistent clustering in *CoCon* compared to *CPC*.



Figure 5: Soft Alignment of videos from UCF101 test split using CoCon pre-trained on UCF101. The first pair of videos involves pull-ups; observe the periodicity captured in the heatmap. The second involves high-jumps; notice that we are roughly able to align the running and jumping phases though they happen at different times. Heatmaps (right) represent relative block similarities from different time-steps of the videos. The color of the frame boxes describe the associated actions; matching colors broadly represent similar action stages.

Effect of action classes on performance Figure 4 shows the classes which observe the least and highest performance improvements when co-trained with multiple views. We observe a loose pattern where action classes involving distinguishable physical movements see larger improvements. We can argue this is because we use views which are suitable for physically intensive actions.

Inter-Class Relationships In order to study consistency of structure across different views, we look at relationships between classes by inferring their similarities through our learned features. We compare cosine similarities across videos from different classes and compute the most similar four classes for each action. We repeat the process for all views and look at the consistency of the results. We only display classes which are amongst the closest ones across all views. Table 4 summarizes our results. We see the detected nearest actions are semantically related to the original actions. In the cases of PlayingCello, we encounter a cluster of categories involving playing instruments. Similarly for Basketball, we can see emergence of sports-based relationships even though there isn't any visual commonality between the categories. It's worth noting that as these nearest classes are consistent across different views, our approach cannot cheat to generate them i.e. it cannot look at 'background crowd' or 'green field' and infer that a video clip is related to sports. Since views such as Optical-Flow, SegMasks and KeypointHeatmap do not have such information and are very low-dimensional.

Action Alignment Even though we only use self-supervision, our embeddings are able to capture relevant

semantics through our multi-view approach allowing loose alignment between videos. To compute this soft alignment, we divide each video into 18 blocks and compute block-level features. We then utilize relative cosine similarities to infer associations between the videos. Figure 5 highlights a few examples. Notice the periodicity implicitly present in some actions (e.g. pullups) captured through the heatmap allowing us to perform non-linear alignment.

5. Conclusion

We propose a cooperative version of contrastive learning, called CoCon, for self-supervised video representation learning. By leveraging relationships across views, we encourage our self-supervised learning objective to be aligned with the underlying semantics. We demonstrate the effectiveness of our approach on the downstream task of action classification, and illustrate the semantic structure of our representation. We show that additional input views generated by off-the-shelf computer vision algorithms can lead to significant improvements, even though they are noisy and derived from an existing modality i.e. RGB. As these views are 'freely' available, this shows the feasibility of utilizing multi-view approaches on datasets which are not traditionally considered multi-view. We hope this enables the ability to leverage multi-view learning algorithms and observe performance gains even on single-view datasets.

Acknowledgement. This work has been supported by Toyota Research Institute (TRI). This article solely reflects the opinions and conclusions of its authors and not TRI or any entity associated with TRI.

References

- [1] Humam Alwassel, Dhruv Mahajan, Bruno Korbar, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering. *arXiv preprint arXiv:1911.12667*, 2019.
- [2] Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019.
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [4] Hanneke EM Den Ouden, Peter Kok, and Floris P De Lange. How prediction errors shape perception, attention, and motivation. *Frontiers in psychology*, 3:548, 2012.
- [5] Carl Doersch and Andrew Zisserman. Multi-task self-supervised visual learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2051–2060, 2017.
- [6] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *European conference on computer vision*, pages 184–199. Springer, 2014.
- [7] Gunnar Farneback. Two-frame motion estimation based on polynomial expansion. In *Scandinavian conference on Image analysis*, pages 363–370. Springer, 2003.
- [8] Chuang Gan, Boqing Gong, Kun Liu, Hao Su, and Leonidas J Guibas. Geometry guided convolutional neural networks for self-supervised video representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5589–5597, 2018.
- [9] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304, 2010.
- [10] Tengda Han, Weidi Xie, and Andrew Zisserman. Video representation learning by dense predictive coding. In *Workshop on Large Scale Holistic Video Understanding, ICCV*, 2019.
- [11] Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning. In *European Conference on Computer Vision*, 2020.
- [12] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *arXiv preprint arXiv:2010.09709*, 2020.
- [13] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6546–6555, 2018.
- [14] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019.
- [15] Jakob Hohwy. *The predictive mind*. Oxford University Press, 2013.
- [16] Anthony Hu, Fergal Cotter, Nikhil Mohan, Corina Gurau, and Alex Kendall. Probabilistic future prediction for video scene understanding. *arXiv preprint arXiv:2003.06409*, 2020.
- [17] Dinesh Jayaraman and Kristen Grauman. Slow and steady feature analysis: higher order temporal coherence in video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3852–3861, 2016.
- [18] Longlong Jing and Yingli Tian. Self-supervised spatiotemporal feature learning by video geometric transformations. *ArXiv*, abs/1811.11387, 2018.
- [19] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [20] Dahun Kim, Donghyeon Cho, and In So Kweon. Self-supervised video representation learning with space-time cubic puzzles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8545–8552, 2019.
- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [22] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Colorization as a proxy task for visual understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6874–6883, 2017.
- [23] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 667–676, 2017.
- [24] Yingming Li, Ming Yang, and Zhongfei Zhang. A survey of multi-view representation learning. *IEEE transactions on knowledge and data engineering*, 31(10):1863–1883, 2018.
- [25] William Lotter, Gabriel Kreiman, and David Cox. Deep predictive coding networks for video prediction and unsupervised learning. *arXiv preprint arXiv:1605.08104*, 2016.
- [26] Zelun Luo, Boya Peng, De-An Huang, Alexandre Alahi, and Li Fei-Fei. Unsupervised learning of long-term motion dynamics for videos. *CoRR*, abs/1701.01821, 2017.
- [27] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015.
- [28] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision*, pages 527–544. Springer, 2016.
- [29] Andriy Mnih and Koray Kavukcuoglu. Learning word embeddings efficiently with noise-contrastive estimation. In *Advances in neural information processing systems*, pages 2265–2273, 2013.
- [30] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pages 69–84. Springer, 2016.
- [31] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [32] 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. *arXiv preprint arXiv:2003.04298*, 2020.

- [33] Zhongzheng Ren and Yong Jae Lee. Cross-domain self-supervised multi-task feature learning using synthetic imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 762–771, 2018.
- [34] Mohammad Sabokrou, Mohammad Khalooei, and Ehsan Adeli. Self-supervised representation learning via neighborhood-relational encoding. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8010–8019, 2019.
- [35] Nawid Sayed, Biagio Brattoli, and Björn Ommer. Cross and learn: Cross-modal self-supervision. *CoRR*, abs/1811.03879, 2018.
- [36] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014.
- [37] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [38] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [39] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhutdinov. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, pages 843–852, 2015.
- [40] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *CoRR*, abs/1906.05849, 2019.
- [41] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Anticipating the future by watching unlabeled video. *CoRR*, abs/1504.08023, 2015.
- [42] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. *CoRR*, abs/1609.02612, 2016.
- [43] Xiaolong Wang and Abhinav Gupta. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2794–2802, 2015.
- [44] Donglai Wei, Joseph J Lim, Andrew Zisserman, and William T Freeman. Learning and using the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8052–8060, 2018.
- [45] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [46] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3733–3742, 2018.
- [47] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016.