

Motion-Focused Contrastive Learning of Video Representations*

Rui Li[†], Yiheng Zhang[‡], Zhaofan Qiu[‡], Ting Yao[‡], Dong Liu[†], and Tao Mei[‡]

[†] University of Science and Technology of China, Hefei, China [‡] JD AI Research, Beijing, China

liruid@mail.ustc.edu.cn, {yihengzhang.chn, zhaofanqiu, tingyao.ustc}@gmail.com

dongeliu@ustc.edu.cn, tmei@live.com

Abstract

Motion, as the most distinct phenomenon in a video to involve the changes over time, has been unique and critical to the development of video representation learning. In this paper, we ask the question: how important is the motion particularly for self-supervised video representation learning. To this end, we compose a duet of exploiting the motion for data augmentation and feature learning in the regime of contrastive learning. Specifically, we present a Motion-focused Contrastive Learning (MCL) method that regards such duet as the foundation. On one hand, MCL capitalizes on optical flow of each frame in a video to temporally and spatially sample the tubelets (i.e., sequences of associated frame patches across time) as data augmentations. On the other hand, MCL further aligns gradient maps of the convolutional layers to optical flow maps from spatial, temporal and spatio-temporal perspectives, in order to ground motion information in feature learning. Extensive experiments conducted on R(2+1)D backbone demonstrate the effectiveness of our MCL. On UCF101, the linear classifier trained on the representations learnt by MCL achieves 81.91% top-1 accuracy, outperforming ImageNet supervised pre-training by 6.78%. On Kinetics-400, MCL achieves 66.62% top-1 accuracy under the linear protocol.

1. Introduction

The key difference between video and image is the dimension of time, which derives a particular form of motion information in a video. The state-of-the-art works often delve into motion in different ways, e.g., long/short term dependencies [29, 43, 21], temporal structure/orders [19, 20, 47], and temporal pooling [42, 49], to enhance video understanding. The underlying foundation behind these advances generally originates from the improvement of representation learning via the exploration of motion information. Most recently, self-supervised representation learning is gaining significant momentum [2, 4, 14], and

the number of self-supervised learning papers practically exploded. In particular, contrastive learning, as a memory-based self-supervised learning approach, is extended to video domain [13, 50] and further closes the gap between self-supervised and supervised video representation learning. A valid question then emerges as how important is the motion for self-supervised video representation learning?

In an effort to answer the question, we look into the problem, in the context of contrastive learning, from two different perspectives: 1) leveraging motion information in achieving data augmentations, and 2) taking motion into account in the optimization of feature learning. In a video, the motion of different regions is inherently various and the velocity of motion measures the rate of change in position of the region with respect to a frame of reference. In general, the regions with larger velocities have much richer information and are potentially more advantageous for contrastive learning. As a result, we capitalize on motion information from both spatial and temporal dimensions to carefully sample the sequences of patches across frames, i.e., tubelets, as augmentations, and examine how the first issue affects self-supervised video representation learning. To study the second one, we encourage the grounding of motion information explicitly in feature learning by aligning gradient maps of the convolutional layers to motion (optical flow) maps. As such, feature learning executes the optimization with respect to motion information.

To consolidate the idea of exploring the motion information in video sequence for self-supervised video representation learning, we present a novel Motion-focused Contrastive Learning (MCL) method. Specifically, we leverage unsupervised TV-L1 algorithm [53] to extract the dense optical flow of each frame in a video and compute the motion boundaries as in [6] to obtain the motion map. A video is divided into a set of fixed-length video clips and the spatio-temporal motion map (ST-motion) of each video clip consists of the sequential motion maps of all the frames in the clip. MCL then performs a motion-focused spatio-temporal sampling to select tubelets as data augmentations. Technically, MCL applies a 3D average pooling on the

*This work was performed at JD AI Research.

spatio-temporal motion map to measure the clip-level motion, which indicates the degree of motion of each clip. The clips with relatively large clip-level motion are chosen as the clip candidates for temporal augmentation. Next, MCL employs a temporal pooling on the spatio-temporal motion map of each clip candidate to estimate the motion from spatial viewpoint (S-motion) and localize the spatial patches, which are temporally consistent across frames, as tubelets. Furthermore, in feature learning, MCL extracts the spatio-temporal motion map of each tubelet and executes a spatial/temporal pooling on such spatio-temporal motion map to output the motion map from temporal/spatial viewpoint (T/S-motion) of the tubelet. The gradient with regard to the feature map of a convolutional layer by back-propagation is produced in spatial, temporal and spatio-temporal manner, respectively, to align with S-motion, T-motion, ST-motion through minimizing the mean squared error in between. MCL integrates the alignments into contrastive learning framework as constraints in addition to InfoNCE loss.

The main contribution of the work is the proposal of leveraging motion information to boost self-supervised video representation learning on the recipe of contrastive learning. This leads to the elegant views of how to effectively sample spatio-temporal augmentations in terms of motion, and how to integrate motion information into the optimization of feature learning, which are problems not yet fully understood. We demonstrate that our self-supervised method MCL surpasses ImageNet supervised pre-training on two video benchmarks and the experiments on two downstream video tasks also validate our MCL.

2. Related work

Unsupervised video representation learning aims to explore the intrinsic properties in unlabeled videos to learn video representation. The research in this direction has proceeded along two different dimensions: transformation-based methods [1, 15, 16], and temporal context-based approaches [8, 11, 12, 19, 47]. Transformation-based methods are optimized to predict the transformation parameters from the transformed videos. Jing *et al.* [15] introduce a pretext task to estimate the rotation angle applied to videos. 3D ST-puzzle [16] proposes a self-supervised task to classify the arrangement of cropped spatio-temporal pieces. SpeedNet [1] learns video representation by estimating the speed or pace of the transformed videos. Temporal context-based approaches focus on exploring the natural temporal relation as supervision. In [19, 24, 47], predicting the order of frames or video clips drives the learning of spatio-temporal representation. The video representation in [11, 12] are learnt by a dense encoding of spatio-temporal blocks to recurrently generate the future representations. Dynamonet [8] directly takes the reconstruction of future frames as the pretext task.

Learning with motion. Motion information, as the rep-

resentation of the changes over time, has been studied by researchers for a long time. For instance, in [22, 38, 54], the optical flow is utilized to propagate the adjacent frame representations. In [31], the famous two-stream architecture is devised by applying two 2D CNN architectures separately to RGB frames and optical flows for action recognition. The idea of two-stream architecture is also explored from the perspective of knowledge distillation [5, 33, 45]. Moreover, Wang *et al.* [40] devise a self-supervised pretext task by estimating the motion in unlabeled videos.

Contrastive learning recently has received intensive research attention due to its promising result on self-supervised visual representation learning. The contrastive loss is devised to return low values for similar pairs and high values for dissimilar pairs, which encourages invariant features on the low dimensional manifold. In an early work [44], the contrastive learning is formulated as instance-level classification, and the previously computed features are stored in a memory bank to acquire more negative samples. Momentum Contrast (MoCo) [14] builds a dynamic memory bank to maintain a large number of negative samples with a moving-averaged encoder. [4, 36] further investigate the importance of data augmentations and non-linear in contrastive learning.

In short, our work in this paper mainly focuses on improving contrastive learning for video representation learning through the involvement of motion information. The most closely related works are the contrastive learning frameworks for video representation learning [13, 26, 41, 50]. SeCo [50] and Pace Prediction [41] combines contrastive loss with order prediction and pace estimation pretext task, respectively. CoCLR [13] exploits the complementary information from optical flow and introduces a co-training scheme to improve the spatio-temporal representation. CVRL [26] studies what makes the good data augmentation for video self-supervised learning. Our method differs from these works in that we exploit the motion information in contrastive learning framework from the perspectives of both data augmentation and representation learning.

3. Motion-Focused Contrastive Learning

The basic idea of Motion-focused Contrastive Learning (MCL) is to facilitate self-supervised video representation learning by the awareness of motion information. MCL first pre-estimates the motion map to measure the pixel-level movement in video data. The regions with larger velocities generally contain richer information (e.g., moving object, action, interaction), and therefore have higher priority in MCL. Specifically, our framework integrates the motion map into contrastive learning from the perspectives of both data augmentation and feature optimization. Figure 1 illustrates the overview of our MCL.




Figure 1. An overview of our **Motion-focused Contrastive Learning (MCL)** framework. MCL facilitates the basic contrastive learning framework by the two novel designs: motion-focused video augmentation and motion-focused feature learning. Specifically, the **motion-focused video augmentation** exploits the pre-estimated motion map to produce the 3D tubelets with rich motion information as the inputs to backbone network. The augmentation process can be divided into two parts, i.e., temporal sampling to filter out the relatively static clips, and spatial cropping to select the spatial regions with high moving velocity. For **motion-focused feature learning**, a new motion alignment loss is devised to enforce the backbone network to focus more on the positions with higher temporal dynamics by aligning the magnitude of gradient and motion map in each position. The motion alignment loss is integrated into the contrastive learning framework as constraints in addition to InfoNCE loss. The whole MCL framework is jointly optimized in an end-to-end manner.

3.1. Motion Estimation

In order to explore the motion information in video sequence, MCL starts by estimating the magnitude of motion in each region. Given a video of N frames with the resolution of $H \times W$, we first extract the **optical flow** from each pair of consecutive frames as $\{(\mathbf{u}_1, \mathbf{v}_1), (\mathbf{u}_2, \mathbf{v}_2), \dots, (\mathbf{u}_N, \mathbf{v}_N)\}$ by TV-L1 [53] algorithm. The flow maps $\mathbf{u}_i, \mathbf{v}_i \in \mathbb{R}^{H \times W}$ are the horizontal and vertical displacements of each pixel between i -th frame and $(i + 1)$ -th frame. For the last frame, we manually set $(\mathbf{u}_N, \mathbf{v}_N) = (\mathbf{u}_{N-1}, \mathbf{v}_{N-1})$. These optical flow maps have been proven effective in capturing temporal dynamics and widely utilized in video classification methods [3, 9, 28, 52]. Nevertheless, in our case of measuring the magnitude of movement, the results by optical flow may suffer from stability problem due to camera motion. For example, with large camera motion, the static objects or background pixels also show high moving velocity in optical flow. Hence, we calculate the **motion boundary** proposed in [6] as $(\frac{\partial \mathbf{u}_i}{\partial x}, \frac{\partial \mathbf{u}_i}{\partial y}, \frac{\partial \mathbf{v}_i}{\partial x}, \frac{\partial \mathbf{v}_i}{\partial y})$, i.e., the x- and y- derivatives of optical flow, to eliminate the effect of camera motion. Finally, we define the **motion map** by accumulating the amplitudes in four motion boundary maps as

$$\mathbf{m}_i = \sqrt{\left(\frac{\partial \mathbf{u}_i}{\partial x}\right)^2 + \left(\frac{\partial \mathbf{u}_i}{\partial y}\right)^2 + \left(\frac{\partial \mathbf{v}_i}{\partial x}\right)^2 + \left(\frac{\partial \mathbf{v}_i}{\partial y}\right)^2}, \quad (1)$$

where $\mathbf{m}_i \in \mathbb{R}^{H \times W}$ measures only the moving velocity in i -th frame and ignores moving orientation. Figure 2 shows the input video and the visualizations of optical flow, motion boundary, and motion map. As illustrated in the figure, the motion map is not influenced by camera motion and shows high responses on the de facto moving objects.

To describe the utilization of motion map more clearly, we pre-define three different types of motion map, i.e., ST-motion, S-motion and T-motion, to measure the moving velocity from different aspects. **ST-motion** stacks the motion map of all frames to produce a 3D volume $\mathbf{m}^{\text{ST}} \in \mathbb{R}^{N \times H \times W}$. **S-motion** and **T-motion** averagely pool the motion map through temporal dimension and spatial dimension, respectively:

$$\begin{aligned} \mathbf{m}^{\text{S}} &= P_t(\mathbf{m}^{\text{ST}}) \in \mathbb{R}^{H \times W}, \\ \mathbf{m}^{\text{T}} &= P_s(\mathbf{m}^{\text{ST}}) \in \mathbb{R}^N, \end{aligned} \quad (2)$$

where $P_t(\cdot)$ and $P_s(\cdot)$ are the pooling operations. Please note that the motion estimation process does not require manual labeling. Hence, these motion maps can be treated as an additional label-free supervision for video data.

3.2. Motion-Focused Video Augmentation

Then, the video augmentation in MCL is guided by the acquired motion map and focuses more on the regions with large movements. The rationale behind is to select the better views for contrastive learning that could boost the generalization capability of the learnt representations. Specifically, in contrastive learning, a single instance is split into two views to learn an embedding where the views are relatively closer than the views from other instances. The self-supervised learning therefore will benefit from the mutual information (MI) between views. As studied in [36], to improve the generalization ability to downstream tasks, the “good” views should contain as much task-relevant information while discarding as much irrelevant information in the input as possible. Unlike the framework in [36] that requires the prior knowledge of downstream tasks to select views, MCL exploits the label-free motion map to seek the

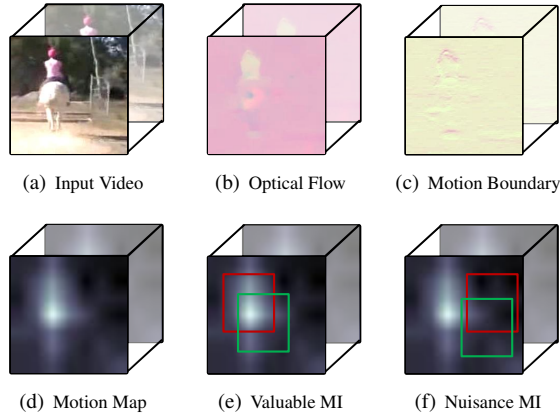


Figure 2. The example of (a) input video sequence, (b) optical flow, (c) motion boundary, (d) motion map, (e) valuable mutual information in the region with large movement and (f) nuisance mutual information in the relatively static region. The red box and green box represent two views generated from the identical instance for contrastive learning.

regions with large movements, which are generally essential for video-related tasks. For example, the views generated in the region with large motion and relatively static region are shown in Figure 2(e) and Figure 2(f), respectively. The views with large movements are more likely to contain the valuable mutual information such as moving object or action, while the views in relatively static region only contain the information in background pixels that are usually considered as nuisance information in contrastive learning.

Here, we introduce a simple way to ensure the framework focusing more on the motion information by only selecting the 3D tubelets with large movement as preprocessing. To achieve this, **motion-focused video augmentation** screens the input video volume by two steps of temporal sampling and spatial cropping, meanwhile maintaining the randomness in the traditional augmentation methods.

Temporal sampling selects the clips (i.e., short frame sequences) with large movements in the video. Taking the T-motion \mathbf{m}^T in Eq.(2) as the frame-level motion map, the movement of each clip is measured by averaging the frame-level movement in \mathbf{m}^T . We take the median across all the candidate clips in an identical video as a threshold, and randomly select one clip with magnitude above the threshold.

Spatial cropping further localizes the cropped spatial region in the sampled clip. It first calculates the S-motion \mathbf{m}^S within the clip. We then take the 90th-percentile in the motion map as a threshold. One bounding box that covers the 80% pixels with higher value than the threshold is randomly selected. Once the region is determined, we crop the same region of each frame in the clip, which is the same as the temporally consistent spatial augmentation in [26].

Please note that this two-step augmentation can also be equivalently implemented by directly seeking the spatio-temporal region with large movement in the ST-motion.

Here we separate it into two steps due to the higher efficiency since temporal sampling only selects a small number of candidate tubelets in the first step. Moreover, the two thresholds in temporal sampling and spatial cropping are relative values within each video/clip. We do not compare the movement magnitude across different videos to avoid over-fitting to the videos with large motion.

Based on the produced clips by temporal sampling and spatial cropping, we follow [26, 50] and also employ color-jittering, random scales, grayscale, blur and mirror.

3.3. Motion-Focused Feature Learning

In addition to video augmentation, we also take the motion map as an additional supervision to enhance the learnt representation. Particularly, we begin by briefly reviewing instance discrimination objective in contrastive learning. Given an encoded query $\mathbf{q} \in \mathbb{R}^d$ and a group of encoded key vectors $\mathcal{K} = \{\mathbf{k}^+, \mathbf{k}_1^-, \mathbf{k}_2^-, \dots, \mathbf{k}_K^-\}$ consisting of one positive key $\mathbf{k}^+ \in \mathbb{R}^d$ and K negative keys $\mathcal{K}^- = \{\mathbf{k}_j^-\}$, where d denotes the embedding dimension. The query and its positive key are usually two views generated from the same instance, while the negative keys are from the other instances. The objective of instance discrimination is to guide the query \mathbf{q} to be similar with the positive key \mathbf{k}^+ while it remains distinct to all negative keys \mathcal{K}^- . Therefore, a prevailing form of contrastive loss is presented in InfoNCE [25] based on a softmax formulation:

$$\mathcal{L}_{\text{NCE}} = -\log \frac{\exp(\mathbf{q}^T \mathbf{k}^+ / \tau)}{\exp(\mathbf{q}^T \mathbf{k}^+ / \tau) + \sum_{i=1}^K \exp(\mathbf{q}^T \mathbf{k}_i^- / \tau)}, \quad (3)$$

where the similarity is measured via dot product, and τ is the temperature hyper-parameter. Here, we follow the recent variant MoCo [14] that proposes the use of a queue to track negative samples from past mini-batches. In this way, all the queuing samples from multiple mini-batches serve as negative keys, and the size of mini-batch can be reduced.

The NCE loss in Eq.(3) performs on the encoded tubelet-level representation, which treats each position in the tubelet equally. As discussed in Section 3.2, the background positions without movements may lead to nuisance information between views. Consequently, we propose a novel **motion alignment loss** (MAL) to explicitly align the feature/gradient maps of convolutional layer and the magnitude in low-level motion map as an additional supervision. Such supervision encourages the network to focus on the positions with large movements in videos. To achieve this, we devise the following variants of motion alignment loss:

(i) *Align feature map*: The simplest way is to align the magnitude of convolutional feature with the motion map. The region with large movement is expected to have high response in the feature map. Formally, let \mathbf{h}_c denote the c -th channel of output feature map from the last convolutional layer of the backbone. The loss function is formulated as the




Figure 3. An illustration of **motion alignment loss** which aligns the gradient maps with ST-motion, S-motion and T-motion.

L2 distance between the summation of \mathbf{h}_c across all channels and ST-motion:

$$\mathcal{L}_{\text{MAL-v1}} = \|\langle \sum_c \mathbf{h}_c \rangle - \langle \mathbf{m}^{\text{ST}} \rangle\|_2^2, \quad (4)$$

where $\langle \cdot \rangle$ is the L2 normalization of the feature/motion map.

(ii) *Align weighted feature map*: Inspired by GradCAM [30], the magnitude of gradient could measure the contribution of each position better. Hence, we calculate the gradient of the similarity between query and its positive key with regard to the convolutional feature as $\mathbf{g}_c = \frac{\partial \mathbf{q}^T \mathbf{k}^+}{\partial \mathbf{h}_c}$. Then, the mean value w_c of the gradient map \mathbf{g}_c can be utilized as the weight of each channel as

$$\mathcal{L}_{\text{MAL-v2}} = \|\langle \text{ReLU}(\sum_c w_c \mathbf{h}_c) \rangle - \langle \mathbf{m}^{\text{ST}} \rangle\|_2^2, \quad (5)$$

where the ReLU operation is added to filter out the regions with negative contribution as in [30].

(iii) *Align weighted gradient map*: We further consider to align the gradient map with motion map. As such, motion information can directly guide the update of representations. Specifically, we replace the weighted feature map in Eq.(5) with weighted gradient map:

$$\mathcal{L}_{\text{MAL-v3}} = \|\langle \text{ReLU}(\sum_c w_c \mathbf{g}_c) \rangle - \langle \mathbf{m}^{\text{ST}} \rangle\|_2^2. \quad (6)$$

The comparisons between different loss functions will be discussed in the experiments, and the alignment of weighted gradient map is used as the default motion alignment loss.

To emphasize the alignment from temporal/spatial aspects, we additionally minimize the distance between spatially/temporally pooled weighted gradient map and T/S-motion, as shown in Figure 3. Specifically, we define the spatio-temporal weighted gradient map as $\bar{\mathbf{g}}^{\text{ST}} = \langle \text{ReLU}(\sum_c w_c \mathbf{g}_c) \rangle$, and further pool such map through temporal/spatial dimension to produce spatial/temporal weighted gradient map $\bar{\mathbf{g}}^{\text{S}}/\bar{\mathbf{g}}^{\text{T}}$, respectively. The three gradient maps are then aligned to the corresponding motion map:

$$\mathcal{L}_{\text{MAL}} = \|\bar{\mathbf{g}}^{\text{S}} - \langle \mathbf{m}^{\text{S}} \rangle\|_2^2 + \|\bar{\mathbf{g}}^{\text{T}} - \langle \mathbf{m}^{\text{T}} \rangle\|_2^2 + \|\bar{\mathbf{g}}^{\text{ST}} - \langle \mathbf{m}^{\text{ST}} \rangle\|_2^2. \quad (7)$$

The overall training objective in our MCL is formulated as a multitask loss by integrating the instance discrimination loss and motion alignment loss, which is written as $\mathcal{L} = \mathcal{L}_{\text{NCE}} + \mathcal{L}_{\text{MAL}}$, where we empirically treat each loss equally and simply summate the two losses.

4. Experiments

We verify the merit of our MCL via various empirical evidences. The evaluation protocols include: 1) linear classification on action recognition, directly trained on frozen pre-learned features, and 2) fine-tuning the networks initialized from pre-learned network for downstream tasks of action recognition and video retrieval.

4.1. Datasets

UCF101 [32] contains about 13K videos from 101 action categories, which are grouped into 9.5K and 3.7K videos for training and testing. HMDB51 [18] consists of 7K videos from 51 action classes, and each split includes 3.5K and 1.5K videos for training and testing. Kinetics-400 dataset [3] is a large-scale action recognition benchmark and contains around 300K videos from 400 action categories. The 300K videos are divided into 240K, 20K and 40K in training, validation and test sets, respectively. Note that the labels of test set are not publicly available and here we report the performances on validation set.

4.2. Experimental Settings

Implementation Details. We exploit the backbones of R(2+1)D based on ResNet-50 [27, 37] or S3D [46] plus an MLP head as video encoder for contrastive training. Note that MLP head only works for training and is disabled on downstream tasks. The inputs of tubelets to the backbone are with the size of $16 \times 224 \times 224$, and each tubelet consists of 16 frames with a temporal stride 2. The MLP head takes 3D global pooling features as the inputs and embeds the feature into $128d$ via two fully-connected layers (2048×2048 and 2048×128). We apply L2 normalization to the output vector from the MLP head, which is then used as the encoded feature of query or keys. In the implementations, we set the size of mini-batch and memory to 64 and 131,072, respectively. The momentum coefficient α is fixed to 0.999 for momentum update of video encoder and the temperature τ in infoNCE loss is 0.1. As implemented in [14], shuffling BN is employed to avoid data leaks. For the optimization of video encoder, we use the momentum SGD with an initial learning rate 0.01 which is annealed down to zero via a cosine decay. Following [50], the network is trained for 400 epochs on Kinetics-400 and 200 epochs on UCF101 with the network initialization by MoCo on ImageNet.

Linear Protocol. We directly utilize the backbone pre-learned by MCL as a feature extractor and examine the quality of frozen features on UCF101 and Kinetics-400 datasets. Specifically, we densely sample 20 clips from each video and resize each clip with the short edge of 256. The resized clip is then cropped to 224×224 via the standard center crop, which is fed into the feature extractor to output the clip-level feature, and the video-level feature is the

Table 1. Performance contribution of each design in MCL with different backbone networks (All models are learnt on UCF101 and evaluated under linear protocol).

TA	SA	\mathcal{L}_t	\mathcal{L}_s	\mathcal{L}_{st}	ResNet-50		Inception V1	
					R2D	R(2+1)D	I2D	S3D
					76.51	77.98	74.39	76.47
					76.47	78.26	74.22	76.79
✓					77.32	79.54	75.02	77.95
✓	✓				77.16	79.64	74.89	78.01
✓	✓	✓			-	80.26	-	78.62
✓	✓		✓		77.84	80.83	75.64	79.28
✓	✓			✓	-	81.47	-	79.40
✓	✓	✓	✓	✓	-	81.91	-	79.88

average of all clip-level features. Finally, we train a linear SVM (UCF101) or a fully-connected layer followed by softmax (Kinetics-400) on the training set, and evaluate the performance on the corresponding validation set. The top-1 classification accuracy is adopted as the metric.

Downstream Task Evaluations. We use the MCL pre-trained backbone as network initialization for two downstream tasks of action recognition and video retrieval in order to examine the transfer capability of the pre-learned structure. For action recognition, we initialize S3D or R(2+1)D network with the backbone self-supervised trained by MCL, and then fine-tune the backbone plus fully-connected layers in standard supervised setting on UCF101 and HMDB51. For video retrieval, we follow the common practice [47], and use the representations of videos from test set to query the k -nearest neighbours in training set. We directly capitalize on the MCL pre-trained backbone without further fine-tuning to extract video representations. We adopt the Recall at k (R@ k) as the performance metric, and a correct retrieval is counted when the top- k nearest neighbours contain at least one video from the category of the query video.

4.3. An Ablation Study of MCL

We first examine how each design in MCL impacts the overall performance. The baseline of data augmentation randomly samples spatial or spatio-temporal patches from the whole video. Temporal Augmentation (TA) solely exploits the clip-level motion to choose the video clips with relatively large motion and randomly localizes the spatial patches in those clips. In contrast, Spatial Augmentation (SA) randomly selects video clips but employs S-motion to locate spatial patches. \mathcal{L}_s , \mathcal{L}_t , and \mathcal{L}_{st} denotes the three items in Eq.(7) and leverages the motion alignment loss from spatial, temporal and spatio-temporal standpoint, respectively. Please note that R2D and I2D are both 2D backbone networks, and the inputs to the two networks are downgraded from 3D tubelets to 2D patches on single frames. As a result, only the grounding on S-motion, i.e., \mathcal{L}_s , is applicable in this case.

Table 1 summarizes the top-1 accuracy on UCF101 dataset under linear protocol by considering different fac-

Table 2. Comparisons of different ways on motion alignment (All models are learnt on UCF101 and evaluated under linear protocol).

Method	Top-1 Acc. (%)
MCL w/o \mathcal{L}_{MAL}	79.64
MCL w/ \mathcal{L}_{MAL-v1}	80.57
MCL w/ \mathcal{L}_{MAL-v2}	80.99
MCL w/ \mathcal{L}_{MAL-v3}	81.91

tors in MCL with various backbone networks. The results consistently indicate that capitalizing on motion to achieve augmentations in SA exhibits performance boost against augmentation by random sampling across 2D and 3D network backbones. Interestingly, TA leads to improvements on 3D network backbones but results in slight drop on 2D network backbones. We speculate that this may be the result of drastic appearance changes in the frames selected by TA and taking the patches from such single frames as inputs to 2D networks may adversely affect feature learning. Involving both TA and SA in MCL contributes a further increase of accuracy with 3D networks of R(2+1)D and S3D. The comparisons also demonstrate the advantages of grounding motion information in feature learning. As aforementioned, only the alignment in spatial manner, i.e., \mathcal{L}_s , fits 2D backbone networks and boosts up the accuracy from 77.16%/74.89% to 77.84%/75.64% on R2D/I2D. Moreover, allowing motion grounding in terms of \mathcal{L}_t , \mathcal{L}_s , and \mathcal{L}_{st} leads to performance gains in the range of 0.62%, 1.19% and 1.83% on R(2+1)D, and 0.61%, 1.27% and 1.39% on S3D. As expected, executing \mathcal{L}_{st} outperforms the use of \mathcal{L}_s or \mathcal{L}_t . By fusing the three losses, the accuracy of MCL finally reaches 81.91% and 79.88% with the backbone of R(2+1)D and S3D.

4.4. Evaluations on Motion Alignment

We then study the effect of three different ways as defined in Eq.(4), Eq.(5) and Eq.(6) for motion alignment in MCL. Table 2 shows the performance comparisons across the three kinds of alignments. As indicated by the results, the use of motion alignment favors the representation learning. That empirically validates the grounding of motion information in MCL. Among the three ways of alignments, MCL w/ \mathcal{L}_{MAL-v3} benefits from the explicit impact on gradient map and leads to a larger performance gain.

4.5. Evaluations on Linear Protocol

Next, we evaluate MCL under linear protocol to verify the representations learnt by MCL. Table 3 details performance comparisons of different representation learning methods on UCF101 and Kinetics-400 datasets. Overall, our MCL leads to a performance boost against all the other baselines on UCF101. In particular, performing classification on the representations pre-learned by MCL with the backbone of S3D and R(2+1)D achieves 79.88% and 81.91%, respectively. Compared to self-supervised method

Table 3. Performance comparisons of video representations pre-learned by different approaches on UCF101 and Kinetics-400 under linear protocol. * represents the results based on the pre-trained models released by original publications on Github. The training time is estimated on one Nvidia P40 GPU. The # of parameters and Flops are counted on an input image/clip with the resolution used in original publications.

Method	Network	Pretrain Dataset	#Param.	FLOPs	Training Time	Eval Dataset	Top1 Acc.(%)
PRP [51]	R(2+1)D	UCF101	14.4M	21.5G	-	UCF101	32.10*
IIC [35]	R3D-10	UCF101	14.4M	19.9G	-	UCF101	35.13*
CCL [17]	R3D-18+1	Kinetics-400	12.1M	16.4G	-	UCF101	52.10
MemDPC [12]	R-2D3D-34	Kinetics-400	32.4M	25.5G	-	UCF101	54.10
TCLR [7]	R3D-18	UCF101	33.0M	32.9G	-	UCF101	67.70
CoCLR [13]	S3D	UCF101	7.9M	12.0G	2.3 days	UCF101	70.20
ImageNet infla.	R(2+1)D	ImageNet	27.3M	20.4G	-	UCF101	75.13
Supervised ImageNet	R-50	ImageNet	23.5M	4.12G	-	UCF101	73.24
SeCo [50]	R-50	UCF101	23.5M	4.12G	0.3 days	UCF101	76.51
SeCo [50]	R(2+1)D	UCF101	27.3M	20.4G	1.4 days	UCF101	77.98
MCL (Ours)	S3D	UCF101	7.9M	18.4G	2.1 days	UCF101	79.88
MCL (Ours)	R(2+1)D	UCF101	27.3M	20.4G	1.5 days	UCF101	81.91
VTHCL [48]	R3D-50	Kinetics-400	31.7M	-	-	Kinetics-400	37.83
SimCLR infla. [26]	R3D-50	Kinetics-400	31.7M	45.8G	-	Kinetics-400	46.80
VINCE [10]	R-50	Kinetics-400	23.5M	4.12G	17.6 days	Kinetics-400	49.10
ImageNet infla. [26]	R3D-50	Kinetics-400	31.7M	45.8G	-	Kinetics-400	53.50
SeCo [50]	R-50	Kinetics-400	23.5M	4.12G	13.0 days	Kinetics-400	61.91
SeCo [50]	R(2+1)D	Kinetics-400	27.3M	20.4G	76.0 days	Kinetics-400	62.50
CVRL [26]	R3D-50	Kinetics-400	31.7M	45.8G	322.6 days	Kinetics-400	66.10
MCL (Ours)	R(2+1)D	Kinetics-400	27.3M	20.4G	76.5 days	Kinetics-400	66.62

CoCLR with S3D backbone, MCL makes the absolute performance improvement by 9.68% based on the same backbone. Moreover, MCL leads the top-1 accuracy by 3.93% over the best competitor SeCo based on the same R(2+1)D backbone. The results empirically verify the idea of leveraging motion in MCL for self-supervised video representation learning. Similar to the observations on UCF101, pre-training MCL on Kinetics-400 dataset outperforms the baselines. MCL with R(2+1)D backbone obtains 66.62% top-1 accuracy and leads to 4.12% performance gain over SeCo with the same backbone. Compared to CVRL which requires $4\times$ the training time and $2\times$ FLOPs of MCL, MCL also leads to an accuracy boost of 0.52%.

4.6. Evaluations on Downstream Tasks

Another common protocol in self-supervised learning is to take the network pre-training as network initialization and fine-tune all layers on downstream tasks. Table 4 shows the comparisons of pre-training the networks by different models and then supervised fine-tuning on UCF101 and HMDB51 for action recognition, which are the most widely-adopted evaluations in the literature. The results across pre-training on UCF101 and Kinetics-400 datasets constantly indicate that our MCL exhibits better performances than all the baselines. Fine-tuning the networks pre-trained on UCF101 by MCL achieves 90.58% and 63.52% on UCF101 and HMDB51, respectively, leading to apparent improvements over VCOP, CoCLR and SeCo. Compared to SeCo which reports the best known results, MCL with R(2+1)D backbone boosts up the accu-

Table 4. Performance comparisons on UCF101 and HMDB51 for downstream task of action recognition.

Method	Network	Pretrain Dataset	UCF101	HMDB51
OPN [19]	VGG	UCF101	59.60	23.80
VCOP [47]	R(2+1)D	UCF101	72.40	30.90
CoCLR [13]	S3D	UCF101	81.40	52.10
BE [39]	R3D-34	UCF101	83.40	53.70
SeCo [50]	R-50	UCF101	83.39	50.19
SeCo [50]	R(2+1)D	UCF101	89.82	56.40
MCL (Ours)	S3D	UCF101	90.58	63.52
MCL (Ours)	R(2+1)D	UCF101	90.40	61.30
3D-RotNet [15]	R3D-18	Kinetics-400	62.90	23.80
ST-Puzzle [16]	R3D-18	Kinetics-400	63.90	33.70
DPC [11]	R-2D3D-34	Kinetics-400	75.70	35.70
MemDPC [12]	R-2D3D-34	Kinetics-400	78.10	41.20
SpeedNet [1]	S3D	Kinetics-400	81.10	48.80
CoCLR [13]	S3D	Kinetics-400	87.90	54.60
BE [39]	R3D-34	Kinetics-400	87.10	56.20
SeCo [50]	R-50	Kinetics-400	88.26	55.55
CVRL [26]	R3D-50	Kinetics-400	92.20	66.70
CBT [34]	S3D	K600+	79.50	44.60
DynamoNet [8]	STCNet	YouTube8M-1	88.10	59.90
MCL (Ours)	R(2+1)D	Kinetics-400	93.41	69.08

racy from 89.82%/56.40% to 90.40%/61.30% on UCF101 and HMDB51 datasets. Pre-training the networks on the larger Kinetics-400 dataset by MCL further improves the accuracy on UCF101 and HMDB51 to 93.41% and 69.08%, and leads the accuracy by 1.21% and 2.38% against CVRL. Notably, MCL is also superior to DynamoNet pre-trained on a subset of YouTube-8M with $2\times$ size of Kinetics-400, which is impressive.

Table 5. Performance comparisons on UCF101 and HMDB51 for downstream task of video retrieval.

Method	Network	Pretrain Dataset	UCF101				HMDB51			
			R @ 1	R @ 5	R @ 10	R @ 20	R @ 1	R @ 5	R @ 10	R @ 20
OPN [19]	VGG	UCF101	19.9	28.7	34.0	40.6	-	-	-	-
VCOP [47]	R3D-18	UCF101	14.1	30.3	40.4	51.1	7.6	22.9	34.4	48.8
VCP [23]	R3D-18	UCF101	18.6	33.6	42.5	53.5	7.6	24.4	36.3	53.6
MemDPC [12]	R-2D3D-18	UCF101	20.2	40.4	52.4	64.7	7.7	25.7	40.6	57.7
SpeedNet [1]	S3D-G	Kinetics-400	13.0	28.1	37.5	49.5	-	-	-	-
PRP [51]	R(2+1)D	UCF101	20.3	34.0	41.9	51.7	8.2	25.3	36.2	51.0
BE [39]	R3D-34	UCF101	-	-	-	-	11.9	31.3	44.5	60.5
CoCLR-RGB [13]	S3D	UCF101	53.3	69.4	76.6	82.0	23.2	43.2	53.5	65.5
CoCLR-2Stream [13]	S3D	UCF101	55.9	70.8	76.9	82.5	26.1	45.8	57.9	69.7
MCL (Ours)	S3D	UCF101	67.0	80.8	86.3	90.8	26.7	52.5	67.0	79.3
MCL (Ours)	R(2+1)D	UCF101	68.6	82.2	87.2	92.0	29.0	55.5	68.9	80.4




Figure 4. From left to right: the first frame from video clip, S-motion, Grad-CAM for SeCo, Grad-CAM for MCL.

Table 5 summarizes the comparisons on UCF101 and HMDB51 for video retrieval task, which is to examine the semantics of k training videos nearest to the query video from test set in the representation space. As indicated by the results across different depths of Recall, MCL yields higher scores than other methods on the two datasets. Taking S3D as the backbone, MCL with only RGB inputs still leads the Recall@1 score by 11.1% and 0.6% over CoCLR with the two-stream inputs of RGB and optical flow modalities on UCF101 and HMDB51. The use of more powerful R(2+1)D backbone in MCL further contributes a Recall@1 increase of 1.6% and 2.3%. The results successfully demonstrate the transferability of the pre-trained structure by MCL to different downstream tasks.

4.7. Visualizing Self-supervised Representation

In order to explain the representations learnt by MCL, we generate the saliency map through GradCAM [30] to indicate the importance of each pixel in presenting the discrimination of the video clip. Figure 4 visualizes the start frame of a video clip, S-motion of the clip, the saliency map produced on the representations learnt by SeCo and our MCL with R(2+1)D backbone. Note that we compute

a 3D saliency map on the video clip and perform temporal pooling to depict the map here. MCL benefits from the leverage of motion and the saliency map on the representations learnt by MCL aligns S-motion more nicely than that by SeCo. More importantly, the regions of high importance effectively provide the visual evidence for describing the video clip and thus the representations learnt by MCL are potentially more robust. That again proves the utilization of motion in MCL for self-supervised representation learning.

5. Conclusion

We have presented a Motion-focused Contrastive Learning (MCL) method, which explores the motion information for improving self-supervised video representation learning. Particularly, we study the problem via leveraging motion to achieve data augmentations and enhance feature learning in contrastive learning framework. To materialize our idea, we extract the motion map of each frame and form a sequence of motion maps in a video clip as the spatio-temporal motion map of the clip. The output value of 3D average pooling on the spatio-temporal motion map is taken as the measure of clip-level motion, which acts as an indicator to select the clip candidates for temporal augmentation. Then, we perform temporal pooling on spatio-temporal motion map of each clip candidate to estimate the motion of every spatial position along the time and localize the spatial patches temporally consistent across frames, as tubelets. MCL employs such tubelets as data augmentations for contrastive learning and further aligns gradients of the convolutional layers to motion maps of the tubelets from spatial, temporal and spatio-temporal aspects. Extensive experiments on UCF101 and Kinetics-400 datasets validate our MCL. More remarkably, self-supervised pre-training MCL is superior to fully-supervised ImageNet pre-training.

Acknowledgments. This work was supported in part by the Natural Science Foundation of China under Grants 62036005 and 62021001, and in part by the Fundamental Research Funds for the Central Universities under contract WK3490000005.

References

- [1] 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 *CVPR*, 2020.
- [2] Qi Cai, Yu Wang, Yingwei Pan, Ting Yao, and Tao Mei. Joint contrastive learning with infinite possibilities. In *NeurIPS*, 2020.
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020.
- [5] Nieves Crasto, Philippe Weinzaepfel, Karteek Alahari, and Cordelia Schmid. Mars: Motion-augmented rgb stream for action recognition. In *CVPR*, 2019.
- [6] Navneet Dalal, Bill Triggs, and Cordelia Schmid. Human detection using oriented histograms of flow and appearance. In *ECCV*, 2006.
- [7] Ishan Dave, Rohit Gupta, Mamshad Nayeem Rizve, and Mubarak Shah. Tclr: Temporal contrastive learning for video representation. *arXiv preprint arXiv:2101.07974*, 2021.
- [8] Ali Diba, Vivek Sharma, Luc Van Gool, and Rainer Stiefelhagen. Dynamonet: Dynamic action and motion network. In *ICCV*, 2019.
- [9] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, 2016.
- [10] Daniel Gordon, Kiana Ehsani, Dieter Fox, and Ali Farhadi. Watching the world go by: Representation learning from unlabeled videos. *arXiv preprint arXiv:2003.07990*, 2020.
- [11] Tengda Han, Weidi Xie, and Andrew Zisserman. Video representation learning by dense predictive coding. In *ICCV Workshop on Large Scale Holistic Video Understanding*, 2019.
- [12] Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning. In *ECCV*, 2020.
- [13] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. In *NeurIPS*, 2020.
- [14] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020.
- [15] Longlong Jing, Xiaodong Yang, Jingen Liu, and Yingli Tian. Self-supervised spatiotemporal feature learning via video rotation prediction. *arXiv preprint arXiv:1811.11387*, 2018.
- [16] Dahun Kim, Donghyeon Cho, and In So Kweon. Self-supervised video representation learning with space-time cubic puzzles. In *AAAI*, 2019.
- [17] Quan Kong, Wenpeng Wei, Ziwei Deng, Tomoaki Yoshinaga, and Tomokazu Murakami. Cycle-contrast for self-supervised video representation learning. In *NeurIPS*, 2020.
- [18] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *ICCV*, 2011.
- [19] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. In *ICCV*, 2017.
- [20] Dong Li, Zhaofan Qiu, Yingwei Pan, Ting Yao, Houqiang Li, and Tao Mei. Representing videos as discriminative sub-graphs for action recognition. In *CVPR*, 2021.
- [21] Dong Li, Ting Yao, Zhaofan Qiu, Houqiang Li, and Tao Mei. Long short-term relation networks for video action detection. In *ACM MM*, 2019.
- [22] Ding Liu, Zhaowen Wang, Yuchen Fan, Xianming Liu, Zhangyang Wang, Shiyu Chang, and Thomas Huang. Robust video super-resolution with learned temporal dynamics. In *ICCV*, 2017.
- [23] Dezhao Luo, Chang Liu, Yu Zhou, Dongbao Yang, Can Ma, Qixiang Ye, and Weiping Wang. Video cloze procedure for self-supervised spatio-temporal learning. In *AAAI*, 2020.
- [24] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *ECCV*, 2016.
- [25] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [26] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning. In *CVPR*, 2021.
- [27] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3d residual networks. In *ICCV*, 2017.
- [28] Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, and Tao Mei. Optimization planning for 3d convnets. In *ICML*, 2021.
- [29] Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xinmei Tian, and Tao Mei. Learning spatio-temporal representation with local and global diffusion. In *CVPR*, 2019.
- [30] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, 2017.
- [31] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014.
- [32] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *CRCV-TR-12-01*, 2012.
- [33] Jonathan Stroud, David Ross, Chen Sun, Jia Deng, and Rahul Sukthankar. D3d: Distilled 3d networks for video action recognition. In *WACV*, 2020.
- [34] Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Learning video representations using contrastive bidirectional transformer. *arXiv preprint arXiv:1906.05743*, 2019.
- [35] Li Tao, Xueting Wang, and Toshihiko Yamasaki. Self-supervised video representation learning using inter-intra contrastive framework. In *ACM MM*, 2020.
- [36] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning. In *NeurIPS*, 2020.

- [37] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018.
- [38] Yi-Hsuan Tsai, Ming-Hsuan Yang, and Michael J Black. Video segmentation via object flow. In *CVPR*, 2016.
- [39] Jinpeng Wang, Yuting Gao, Ke Li, Yiqi Lin, Andy J Ma, Hao Cheng, Pai Peng, Feiyue Huang, Rongrong Ji, and Xing Sun. Removing the background by adding the background: Towards background robust self-supervised video representation learning. In *CVPR*, 2021.
- [40] Jiangliu Wang, Jianbo Jiao, Linchao Bao, Shengfeng He, Yunhui Liu, and Wei Liu. Self-supervised spatio-temporal representation learning for videos by predicting motion and appearance statistics. In *CVPR*, 2019.
- [41] Jiangliu Wang, Jianbo Jiao, and Yun-Hui Liu. Self-supervised video representation learning by pace prediction. In *ECCV*, 2020.
- [42] Peng Wang, Yuanzhouhan Cao, Chunhua Shen, Lingqiao Liu, and Heng Tao Shen. Temporal pyramid pooling-based convolutional neural network for action recognition. *IEEE Trans. on CSVT*, 2016.
- [43] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [44] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018.
- [45] Zeyu Xiao, Xueyang Fu, Jie Huang, Zhen Cheng, and Zhiwei Xiong. Space-time distillation for video super-resolution. In *CVPR*, 2021.
- [46] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, 2018.
- [47] Dejing Xu, Jun Xiao, Zhou Zhao, Jian Shao, Di Xie, and Yueting Zhuang. Self-supervised spatiotemporal learning via video clip order prediction. In *CVPR*, 2019.
- [48] Ceyuan Yang, Yinghao Xu, Bo Dai, and Bolei Zhou. Video representation learning with visual tempo consistency. *arXiv preprint arXiv:2006.15489*, 2020.
- [49] Ceyuan Yang, Yinghao Xu, Jianping Shi, Bo Dai, and Bolei Zhou. Temporal pyramid network for action recognition. In *CVPR*, 2020.
- [50] Ting Yao, Yiheng Zhang, Zhaofan Qiu, Yingwei Pan, and Tao Mei. Seco: Exploring sequence supervision for unsupervised representation learning. In *AAAI*, 2021.
- [51] Yuan Yao, Chang Liu, Dezhao Luo, Yu Zhou, and Qixiang Ye. Video playback rate perception for self-supervised spatio-temporal representation learning. In *CVPR*, 2020.
- [52] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, 2015.
- [53] Christopher Zach, Thomas Pock, and Horst Bischof. A duality based approach for realtime tv-l1 optical flow. In *Joint pattern recognition symposium*, 2007.
- [54] Xizhou Zhu, Yujie Wang, Jifeng Dai, Lu Yuan, and Yichen Wei. Flow-guided feature aggregation for video object detection. In *ICCV*, 2017.