

Hierarchical Contrastive Motion Learning for Video Action Recognition

Xitong Yang¹

xyang35@cs.umd.edu

Xiaodong Yang²

<https://xiaodongyang.org>

Sifei Liu²

<https://www.sifeiliu.net>

Deqing Sun³

<https://deqings.github.io>

Larry Davis¹

<http://users.umiacs.umd.edu/~lsd>

Jan Kautz²

<https://jankautz.com>

¹ University of Maryland, College Park

² NVIDIA

³ Google

Abstract

One central question for video action recognition is how to model motion. In this paper, we present hierarchical contrastive motion learning, a novel self-supervised learning framework to extract effective motion representations from raw video frames. Our approach progressively learns a hierarchy of motion features that correspond to different abstraction levels in a network. At each level, an explicit motion self-supervision is provided via contrastive learning to enforce the motion features to capture semantic dynamics and evolve discriminatively for video action recognition. This hierarchical design bridges the semantic gap between low-level movement cues and high-level recognition tasks, and promotes the fusion of appearance and motion information at multiple levels. Our motion learning module is lightweight and flexible to be embedded into various backbone networks. Extensive experiments on four benchmarks show that our approach compares favorably against the state-of-the-art methods yet without requiring optical flow or supervised pre-training.

1 Introduction

Motion provides abundant and powerful cues for understanding the dynamic visual world. A broad range of video understanding tasks benefit from the introduction of motion information, such as action recognition [41, 48], activity detection [53, 56], object tracking [50, 52], etc. Thus, how to extract and model motion is one of the fundamental problems in video understanding. While early methods in this field mostly rely on the pre-computed motion features such as optical flow, recent works have been actively exploiting convolutional neural networks (CNNs) for more effective motion learning from raw video frames [54, 48], encouraged by the success of end-to-end learning in various vision tasks.

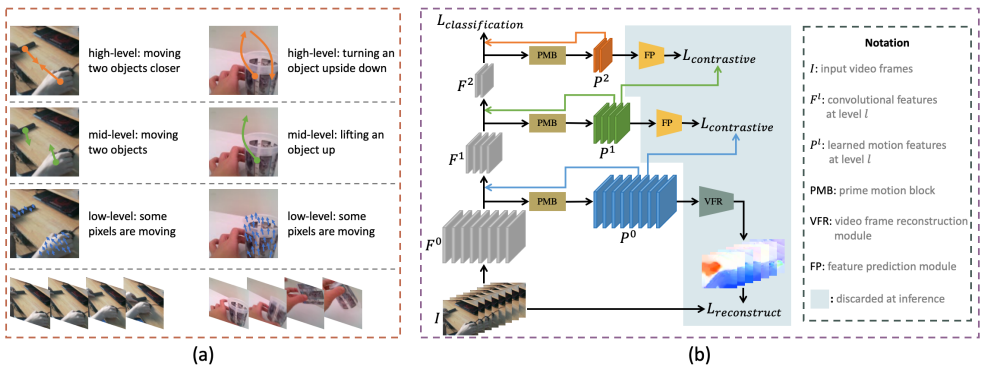


Figure 1: (a) Illustration of the hierarchy of progressively learned motion: from pixel-level and short-range movement to semantic temporal dynamics. (b) Overview of the architecture of the proposed hierarchical contrastive motion learning framework.

A key challenge for end-to-end motion representation learning is to design an effective supervision. Unlike many other tasks that afford plenty of well-defined annotations, the “ground truth motion” is often unavailable or even undefined for motion learning in practice. One popular idea is to extract motion features by means of the action recognition supervision. However, the classification loss is shown to be sub-optimal in this task as it only provides an implied supervision to guide motion learning [43]. This supervision is also prone to be biased towards appearance information since some video action benchmarks can be mostly solved by considering static images without temporal modeling [40]. Recently, some efforts have been made to explore pretext tasks with direct supervisions for motion learning, such as optical flow prediction [64] and video frame reconstruction [59]. Although having shown promising results, such supervisions are restricted to pixel-wise and short-term motion as they hinge on pixel photometric loss and movement between adjacent frames.

In light of the above observations, we introduce a novel self-supervised learning framework that enables explicit motion supervision at multiple feature abstraction levels, which we term hierarchical contrastive motion learning. Specifically, given preliminary motion cues as a bootstrap, our approach progressively learns a hierarchy of motion features in a bottom-up manner, as illustrated in Figure 1(a). This hierarchical design is proposed to bridge the semantic gap between the low-level preliminary motion and the high-level recognition task—analogueous to the findings in neuroscience that humans perceive motion patterns in a hierarchical way [0, 2]. At each level, a discriminative contrastive loss [6, 15] provides an explicit self-supervision to enforce the motion features at current level to predict the future ones at previous level. In contrast to the previous pretext tasks that focus on low-level image details, the contrastive learning encourages the model to learn useful semantic dynamics from previously learned motion features at a lower level, and is more favorable for motion learning at higher levels where the spatial and temporal resolutions of feature maps are low [4, 16, 36]. To acquire the preliminary motion cues to initialize the hierarchical motion learning, we exploit the video frame reconstruction [22] as an auxiliary task such that the whole motion representation learning enjoys a unified self-supervised setup.

We realize the proposed motion learning module via a side network branch, which is lightweight and flexible to be embedded into a variety of backbone CNNs. As shown in Figure 1(b), the side branch (i.e., the shaded region) for self-supervised motion learning is

discarded after training, and only the learned motion features are retained in the form of residual connections [20]. Our hierarchical design also promotes the fusion of appearance and motion by integrating the motion features into the backbone network at multiple abstraction levels. This multi-level fusion paradigm is unachievable for previous motion learning methods [0, 59] that depend solely on low-level motion supervisions.

We summarize our main contributions as follows. (1) We propose a new learning framework for motion representation learning from raw video frames. (2) We advance contrastive learning to a hierarchical design that bridges the semantic gap between low-level motion cues and high-level recognition tasks. (3) To our knowledge, this work provides the first attempt to empower contrastive learning in motion representation learning for large-scale video action recognition. (4) Our approach achieves superior results on four benchmarks without relying on off-the-shelf motion features or supervised pre-training.

2 Related Work

Action Recognition and Motion Extraction. A large family of video action research focuses on motion modeling [19, 33, 41, 52, 54, 55]. For example, the two-stream networks model temporal information by leveraging external motion inputs such as optical flow and frame differences [33, 41]. 3D CNNs [19, 52] and RNNs [54, 55] are also widely used to simultaneously model appearance and motion.

To learn more explicit motion representations, ActionFlowNet [34] uses pre-computed optical flow as an additional supervision to encode motion together with appearance. However, the approach requires the extraction of optical flow which is time-consuming and infeasible for large-scale datasets. Recently, TVNet [9] formulates the TV-L1 algorithm in a customized network layer, producing optical flow like motion features to complement static appearance. A differentiable representation flow layer is also developed in [57]. Inspired by the correlation layer in FlowNet [8], CorrNet [48] and MSNet [25] leverage the correlation operation to extract motion from convolutional features. All these studies employ classification loss as an indirect supervision or mimic optical flow design to learn motion extraction, while our approach explicitly realizes motion learning through the proposed hierarchical contrastive learning in a fully self-supervised way.

Self-Supervised Learning (SSL). To take advantage of the abundant unlabeled videos, numerous methods have been developed for sequence learning by various self-supervisory signals, such as frame interpolation [0, 35], sequence ordering [10, 26], speed prediction [0, 49], future prediction [29, 58], and cross-sensor motion regularization [50]. While these methods do not require pre-trained networks nor video annotations, they tend to focus on the low-level information and may not effectively capture the high-level temporal dynamics.

Recently, contrastive learning has been explored for self-supervised learning of video representations [16, 18, 38, 46, 53]. Different pretext tasks are designed to facilitate the learning of spatio-temporal representations, such as future prediction [16, 17], augmentation invariance [53], and multi-view co-training [18]. Unlike these existing works that leverage SSL to learn discriminative video representations, our work focuses on **extracting motion information** from raw video frames in a self-supervised manner. Our motion learning can be naturally integrated into the supervised training process and further boosts the supervised learning performance on large-scale datasets. This is not achievable for previous SSL methods because they only serve as a network pre-training process.

3 Approach

Given the convolutional features $\{\mathcal{F}^0, \dots, \mathcal{F}^{L-1}\}$ at different levels of a backbone network, our aim is to learn a **hierarchy** of motion representations $\{\mathcal{P}^0, \dots, \mathcal{P}^{L-1}\}$ for the corresponding levels. Here, L indicates the total number of abstraction levels. As the first step, we employ video frame reconstruction as the supervision to obtain the preliminary motion cues \mathcal{P}^0 , which function as a bootstrap for the following hierarchical motion learning. With that, we progressively learn the motion features in a bottom-up manner. At each level $l > 0$, we learn the motion features \mathcal{P}^l by enforcing them to predict the future motion features at the previous level $l - 1$, as described in Sec. 3.1. We use the contrastive loss as an objective such that \mathcal{P}^l is trained to capture semantic temporal dynamics from \mathcal{P}^{l-1} . The learned motion features at each level are integrated into a backbone network via residual connections to perform appearance and motion feature fusion: $\mathcal{Z}^l = \mathcal{F}^l + g^l(\mathcal{P}^l)$, where $g^l(\cdot)$ is used to match the feature dimensions. After learning motion at all levels, we jointly train the whole network for action recognition in a multi-tasking manner, as presented in Sec. 3.2.

3.1 Self-Supervised Motion Learning

Prime Motion Block. We first introduce a lightweight prime motion block (PMB) to transform the convolutional features of a backbone network to more discriminative representations for motion learning. The key component of this block is a cost volume layer, which is inspired by the success of using cost volumes in stereo matching [21] and optical flow estimation [45]. An illustration of PMB is shown in Figure 2 (a).

Given a sequence of convolutional features $\mathcal{F} = \{F_0, \dots, F_{T-1}\}$ with length T , we first conduct a $1 \times 1 \times 1$ convolution to reduce the input channels by $1/\beta$, denoted as $\tilde{\mathcal{F}}$. This operation significantly reduces the computational overhead of prime motion block, and provides more compact representations to reserve the essential information to compute cost volumes. The adjacent features are then re-organized to feature pairs $\tilde{\mathcal{F}}^* = \{(\tilde{F}_0, \tilde{F}_1), \dots, (\tilde{F}_{T-2}, \tilde{F}_{T-1}), (\tilde{F}_{T-1}, \tilde{F}_{T-1})\}$ for constructing the cost volumes. The matching cost between two features is defined as: $\text{cv}_i(x_1, y_1, x_2, y_2) = \text{sim}(\tilde{F}_i(x_1, y_1), \tilde{F}_{i+1}(x_2, y_2))$, where $\tilde{F}_i(x, y)$ denotes the feature vector at time t and position (x, y) , and the cosine distance is used as the similarity function: $\text{sim}(u, v) = u^T v / \|u\| \|v\|$. We replicate the last feature map \tilde{F}_{T-1} to compute their cost volume in order to keep the original temporal resolution. We limit the search range with the max displacement of (x_2, y_2) to be d and use a striding factor of s to handle large displacements without increasing the computation. As a result, the cost volume layer outputs a feature tensor of size $M \times H \times W$, where $M = (2 \times \lfloor d/s \rfloor + 1)^2$ and H, W denote the height and width of a feature map. Note that computing cost volumes is lightweight as it has no learnable parameters and much fewer FLOPs than 3D convolutions. Finally, we combine cost volumes with features obtained after dimension reduction, motivated by the observation that the two features are complementary for localizing motion boundaries.

Preliminary Motion Cues. Although the prime motion block extracts rough motion features from convolutional features, we find that such features are easily biased towards appearance when jointly trained with the backbone network. Thus, an **explicit** motion supervision is vital for more effective motion learning at each level.

To initialize the progressive training, the preliminary motion cues, i.e., \mathcal{P}^0 , are required as a bootstrap. They should encode some low-level but valid movement information to facilitate the following motion learning. Therefore, we adopt video frame reconstruction to

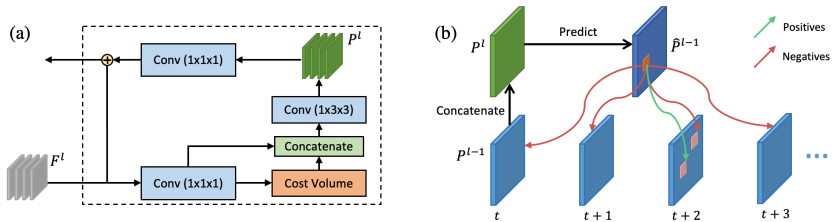


Figure 2: (a) Architecture of the prime motion block. (b) Overview of our contrastive motion learning that adopts the higher-level motion features to predict the future ones at the lower level, as well as the illustration of positive and negative samples.

guide the extraction of preliminary motion cues. This task can be formulated as a self-supervised optical flow estimation problem [22], aiming to produce optical flow to allow frame reconstruction from neighboring frames. Motivated by the success of recent work on estimating optical flow with CNNs [44], we build a simple optical flow estimation module using 5 convolutional layers with dense connections. We make use of the estimated optical flow to warp video frames through bilinear interpolation. The loss function consists of a photometric term that measures the error between the warped frame and the target frame, and a smoothness term that handles the aperture problem that causes ambiguity in motion estimation: $\mathcal{L}_{\text{reconstruct}} = \mathcal{L}_{\text{photometric}} + \zeta \mathcal{L}_{\text{smoothness}}$. We define the photometric error as:

$$\mathcal{L}_{\text{photometric}} = \frac{1}{HWT} \sum_{t=1}^T \sum_{x=1}^W \sum_{y=1}^H \mathbb{1}_{\rho} (I_t(x, y) - \hat{I}_t(x, y)), \quad (1)$$

where \hat{I}_t indicates the warped frame at time t and $\rho(z) = (z^2 + \varepsilon^2)^\alpha$ is the generalized Charbonnier penalty function with $\alpha = 0.45$ and $\varepsilon = 1e^{-3}$, and the indicator function $\mathbb{1} \in \{0, 1\}$ excludes invalid pixels that move out of the image boundary. Additionally, we compute the smoothness term as:

$$\mathcal{L}_{\text{smoothness}} = \frac{1}{T} \sum_{t=1}^T \rho(\nabla_x U_t) + \rho(\nabla_y U_t) + \rho(\nabla_x V_t) + \rho(\nabla_y V_t), \quad (2)$$

where $\nabla_x U/V$ and $\nabla_y U/V$ are the gradients of estimated flow fields U/V in x/y directions.

Hierarchical Motion Learning. Given above preliminary motion cues as a bootstrap, we propose to learn higher-level motion representations using a multi-level self-supervised objective based on the contrastive loss [6, 14, 15]. Our goal is to employ the higher-level motion features as a conditional input to guide the prediction of the future lower-level motion features that are well-learned from a previous step. By this way, the higher-level features are forced to understand a more abstract trajectory that summarizes motion dynamics from the lower-level ones. This objective therefore allows us to extract slowly varying features that progressively correspond to high-level semantic concepts [16, 36].

Formally, let us denote the motion features generated by the prime motion block at level $l > 0$ as $\mathcal{P}^l = \{P_0^l, \dots, P_{T-1}^l\}$. In order to train \mathcal{P}^l , we enforce P_t^l to predict the future motion features at the previous level (i.e., $P_{t+\delta}^{l-1}$), conditioned on the motion feature at the current time P_t^{l-1} , as illustrated in Figure 2(b). In practice, a predictive function f_δ is applied for the motion feature prediction at time step $t + \delta$: $\hat{P}_{t+\delta}^{l-1} = f_\delta([P_t^l, P_t^{l-1}])$, where $[\cdot, \cdot]$ denotes channel-wise concatenation. We employ a multi-layer perception (MLP) with one hidden

layer as the prediction function: $f_{\delta}(x) = W_{\delta}^{(2)} \sigma(W^{(1)}x)$, where σ is ReLU and $W^{(1)}$ is shared across all prediction steps in order to leverage their common information. We define the objective function of each level as a contrastive loss, which encourages the predicted \hat{P}^{l-1} to be close to the ground truth P^{l-1} while being far away from the negative samples:

$$\mathcal{L}_{\text{contrastive}}^l = - \sum_{i \in \mathcal{S}} \left[\log \frac{\exp(\text{sim}(\hat{P}_i^{l-1}, P_i^{l-1})/\tau)}{\sum_{j \in \mathcal{S}} \exp(\text{sim}(\hat{P}_i^{l-1}, P_j^{l-1})/\tau)} \right], \quad (3)$$

where the similarity function is the same cosine similarity used in computing cost volumes, and \mathcal{S} denotes the sampling space of positive and negative samples.

As shown in Figure 2(b), the positive sample of the predicted feature is the ground-truth feature that corresponds to the same video and locates at the same position in both space and time as the predicted one. Similar to [16], we define three types of negative samples for all prediction and ground-truth pairs: spatial negatives, temporal negatives and easy negatives. Considering efficiency, we randomly sample N spatial locations for each video within a mini-batch to compute the loss. Please see the supplementary material for more sampling details. As illustrated in Figure 1(b), the contrastive motion learning is performed for multiple levels until the motion hierarchy of the whole network is built up.

Progressive Training. Training the multi-level self-supervised learning framework simultaneously from the beginning is infeasible, as the lower-level motion features are initially not well-learned and the higher-level prediction would be arbitrary. To facilitate the optimization process, we propose a progressive training strategy that learns motion features for one level at a time, propagating from low-level to high-level. In practice, after the convergence of training at level $l-1$, we freeze all network parameters up to level $l-1$ (therefore fixing the motion features \mathcal{P}^{l-1}), and then start the training for level l . In this way, the higher-level motion features can be stably trained with the well-learned lower-level ones.

3.2 Joint Training for Action Recognition

Our ultimate goal is to improve video action recognition with the learned hierarchical motion features. To integrate the learned motion features into a backbone network, we wrap our prime motion block into a residual block: $\mathcal{Z}^l = \mathcal{F}^l + g^l(\mathcal{P}^l)$, where \mathcal{F}^l is the convolutional features at level l , \mathcal{P}^l is the corresponding motion features obtained in Sec. 3.1, and $g^l(\cdot)$ is a $1 \times 1 \times 1$ convolution. This seamless integration enables end-to-end fusion of appearance and motion information over multiple levels throughout a single unified network, instead of learning them disjointly like two-stream networks [16]. After the motion representations are self-supervised learned at all levels, we add in the classification loss to jointly optimize the total objective, which is a weighted sum of the following losses:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{classification}} + \lambda \mathcal{L}_{\text{reconstruct}} + \sum_l \gamma^l \mathcal{L}_{\text{contrastive}}^l, \quad (4)$$

where λ and γ^l are the weights to balance related loss terms. As illustrated in Figure 1(b), our multi-level self-supervised learning is performed through a side network branch, which can be flexibly embedded into standard CNNs. Furthermore, this self-supervised learning side branch is discarded after training so that our final network can well maintain the efficiency at runtime for inference.

Input	Supervision			Score
	Action	Reconstruction	Contrastive	
Level 1	✓			2.3
		✓		3.1
			✓	–
Level 2	✓			1.7
		✓		2.5
			✓	3.0
Level 3	✓			2.4
		✓		1.7
			✓	3.0

Table 1: Comparison of efficacy scores of the motion features learned at different levels under different supervisory forms.

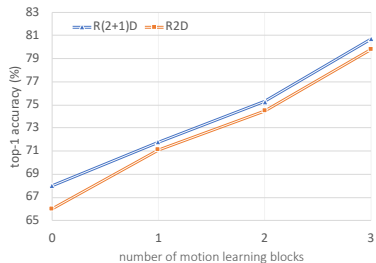


Figure 3: Comparison of the top-1 accuracy on UCF-101 with incrementally adding the proposed motion learning blocks.

4 Experiments

We evaluate the proposed approach on four benchmarks: Kinetics-400 (K400) [3], Something-Something V1&V2 (SS-V1&V2) [13] and UCF-101 [42]. Our motion learning module is generic and can be instantiated with various video networks [8, 47, 51]. In our experiments, we use the standard networks R2D [51] and R(2+1)D [47] as our backbones. We follow the standard recipe in [40] for model training. Note that all models are trained from scratch or self-supervised pre-trained without additional annotations or pre-computed optical flow. More details on datasets and implementations are available in the supplementary material.

4.1 Ablation Study

Supervision for Motion Learning. We first compare the motion features learned at different levels by different supervisions. As our motion learning is based on the self-supervisions that are not directly related with the final action recognition, we first define a measurement to reflect the efficacy of the learned motion features. Towards this goal, we take the extracted motion feature as input and train a lightweight classifier for action recognition on UCF-101. We define the efficacy score as: $\text{Score} = \text{Acc}_{\text{train}} / (\text{Acc}_{\text{train}} - \text{Acc}_{\text{test}})$, where $\text{Acc}_{\text{train}}$ and Acc_{test} indicate the top-1 accuracy on the training and test sets. Intuitively, a higher score implies that the representation is more discriminative (with higher training accuracy) and generalizes better (with a lower performance gap between training and testing).

Table 1 shows the efficacy scores of motion features at different levels with different supervisions, where “action” indicates the supervision by action classification, and “reconstruction” and “contrastive” refer to the supervisions by frame reconstruction and contrastive learning. Levels 1, 2 and 3 correspond to the motion features extracted after res_2 , res_3 and res_4 of R2D. We observe that the self-supervision of low-level frame reconstruction is particularly effective at level 1, but its performance degrades dramatically at higher levels due to lower spatial/temporal resolutions and higher abstraction of convolutional features. In contrast, the proposed self-supervision by hierarchical contrastive learning is more stable over different levels and more effective to model motion dynamics. It is also observed that the self-supervision, with correct choices at different levels, consistently outperforms the supervision by action classification, which is consistent with the findings in [0, 34, 43]. In Figure 4, we visualize the estimated optical flow, the by-product of frame reconstruction at each level, and find that more accurate optical flow indeed presents at lower levels.

Methods	PMB	Self-Sup	FLOPs	UCF-101	SS-V1	K400
Baseline: R2D			1.00×	66.0 / 86.0	36.1 / 68.1	64.8 / 85.1
Ours: R2D	✓		1.18×	71.6 / 89.7	43.6 / 74.7	65.6 / 85.5
Ours: R2D	✓	✓	1.18×	79.8 / 94.4	44.3 / 75.8	67.3 / 86.4
Baseline: R(2+1)D			1.00×	68.0 / 88.2	48.5 / 78.1	66.8 / 86.6
Ours: R(2+1)D	✓		1.11×	73.4 / 92.1	49.2 / 77.9	67.4 / 86.9
Ours: R(2+1)D	✓	✓	1.11×	80.7 / 95.6	50.4 / 78.9	68.3 / 87.4

Table 2: Ablation study on the prime motion block (PMB) and self-supervision (Self-Sup) for action recognition. We report the computational cost and top-1 / top-5 accuracy (%) on the three benchmarks. Models are evaluated using a single clip per video to eliminate the impact of test-time augmentation.




Figure 4: Visualization of the estimated optical flow at different feature abstraction levels. For each group, columns 1-2 are adjacent frames; column 3 is the reference optical flow extracted by [28]; columns 4-6 are the estimated optical flow at levels 1, 2 and 3.

Contributions of Individual Components. We verify the contributions of the proposed components in Table 2. It is obvious that our approach consistently and significantly improves the action recognition accuracy for both 2D and 3D action networks. Our prime motion block provides complementary motion features at multiple levels, and the self-supervision further enhances the representations to encode semantic dynamics. In particular, for the dataset that heavily depends on temporal information like SS-V1, our approach remarkably improves the performance of baseline R2D by 8.2%. For the dataset that is small-scale and tends to overfit to the appearance information like UCF-101, our method improves model generalization and achieves 13.8% improvement. Moreover, our motion learning module only introduces a small overhead to FLOPs of the backbone network.

We next validate the contribution of our motion learning at each level by incrementally adding the proposed motion feature learning block to the baseline. Figure 3 demonstrates the results based on the backbones of R2D and R(2+1)D on UCF-101. We observe that notable gains can be obtained at multiple levels, and the performance gain does not vanish with the increase of motion learning blocks, suggesting the importance of leveraging hierarchical motion information across all levels.

4.2 Comparison with State-of-the-Art

Supervised Action Recognition. We compare our approach with the state-of-the-art methods on the four action recognition benchmarks. Table 3 shows the comparisons on K400 and SS-V1&V2. Without using optical flow or supervised pre-training, our model based on backbone R(2+1)D-101 achieves the best results among the single-stream methods over all three datasets. Our approach also outperforms most two-stream methods, apart from the recent two-stream TSM [27] on SS-V2. As for the datasets that focus more on temporal modeling like SS-V1&V2, 2D networks are usually not able to achieve as good results as 3D models. However, by equipping with the proposed motion learning module, we find that our

Methods	Pre-train	Flow	Arch. (Depth)	GFLOPs×Views [†]	K400	SS-V1	SS-V2
I3D [8]	IN-1K	✓	I3D	216 × N/A	75.7	–	–
S3D-G [20]	IN-1K	✓	S3D	143 × N/A	77.2	48.2	–
Non-local [20]	IN-1K	–	NL R3D (50)	282 × 30	76.5	44.4	–
R(2+1)D [20]	Sports1M	✓	R(2+1)D (34)	304 × 115	75.4	–	–
TSM [20]	IN-1K	–	R2D (50)	65 × 30	74.7	49.7	63.4
ECO [20]	–	✓	ECO _{En}	N/A	70.0	49.5	–
SlowFast [20]	–	–	SlowFast (101)	106 × 30	77.9	–	–
Disentangling [20]	IN-1K	–	Disentangling	N/A	71.5	–	–
D3D [20]	IN-1K	–	S3D	N/A	75.9	–	–
STM [20]	IN-1K	–	R2D (50)	66.5 × 30	73.7	50.7	64.2
Rep. Flow [20]	–	–	R(2+1)D (50)	N/A	77.1	–	–
MARS [8]	–	–	3D ResNeXt (101)	N/A	72.7	–	–
DynamoNet [8]	–	–	STCnet (101)	N/A	77.9	–	–
Ours	–	–	R2D (50)	49 × 30	74.8	46.2	59.4
Ours	–	–	R(2+1)D (101)	150 × 30	78.3	52.8	64.4

Table 3: Comparison of the top-1 accuracy (%) with the state-of-the-art methods on Kinetics-400 (K400) and Something-Something V1&V2 (SS-V1&V2). IN-1K indicates the ImageNet-1K dataset. [†] GFLOPs and views are reported for the experiments on K400.

Methods	Dataset	Acc.
TSN [8]	IN-1K + K400	91.7
TSM [20]	IN-1K + K400	95.9
I3D [8]	IN-1K + K400	95.4
S3D-G [20]	IN-1K + K400	96.8
LGD-3D [20]	IN-1K + K600	97.0
Disentangling [20]	IN-1K + K400	95.9
D3D [20]	IN-1K + K400	97.0
STM [20]	IN-1K + K400	96.2
DynamoNet [8]	YT-8M + K400	97.8
R(2+1)D [20]	K400	96.8
MARS [8]	K400	97.0
Ours , R(2+1)D-101	K400	97.8

Method	Dataset	Res.	Arch. (depth)	Acc.
OPN [20]	UCF	224	VGG (14)	59.8
ActionFlowNet [†] [20]	UCF	224	R3D (18)	83.9
DynamoNet [8]	YT-8M	112	STCNet (133)	88.1
3D-Puzzle [20]	K400	224	R3D (17)	63.9
DPC [20]	K400	128	R-2D3D (33)	75.7
SpeedNet [8]	K400	224	S3D-G (23)	81.1
PacePred [20]	K400	112	R(2+1)D (18)	77.1
MemDPC-RGB [20]	K400	224	R-2D3D (33)	78.1
CVRL [20]	K400	224	R3D (50)	92.1
CoCLR-RGB [†] [20]	K400	128	S3D (23)	87.9
Ours (random init.)	–	128	S3D (23)	77.0
Ours	K400	128	S3D (23)	87.1

Table 4: Comparison with **supervised** learning methods on UCF-101 (3 splits). Models are first supervisedly pre-trained on large-scale datasets and then finetuned on UCF-101.

Table 5: Comparison with the **self-supervised** methods on UCF-101 (split-1). Duration of the pre-training datasets: UCF (1d), K400 (28d), YT-8M (1y). [†] denotes the methods requiring pre-extracted optical flow for model training.

method based on backbone R2D-50 outperforms some 3D models, such as R(2+1)D and NL I3D. Our approach also achieves superior results compared with the most recent work that are specifically designed for temporal motion modeling (i.e., the second group in Table 3). More importantly, our motion learning is fully self-supervised from raw video frames without any supervisions from optical flow or pre-trained temporal stream.

We also conduct experiments on UCF-101 to demonstrate the transferability of our learned features to small-scale datasets. Following the standard setting in previous work [8, 69], we pre-trained our models on K400 for action classification and then finetune the weights on UCF-101. We report the average accuracy over all 3 splits. As shown in Table 4, our approach achieves 97.8% top-1 accuracy, comparable with the state-of-the-art results among the single-stream methods on UCF-101.

Self-Supervised Pre-Training. Self-supervised learning of video representations has been gaining increasing attention in recent years [2, 11, 16, 24, 26]. In addition to hierarchical




Figure 5: Examples of the retrieved videos to reflect the different motion semantics learned by the motion features at different levels.

Motion	Top-2 Classes with Largest Relative Gains
Low-level	“turning the camera left” (73.8) “turning the camera right” (71.7)
Mid-level	“showing smth to the camera” (0 → 19.4) “showing smth behind smth” (0 → 12.2)
High-level	“pulling two ends of smth” (0 → 14.0) “tipping smth with smth in it” (0 → 12.5)

Table 6: Comparison of video classification accuracy using the motion features learned at different levels.

motion learning, our approach can also serve as pre-training of a network. As an example, after self-supervised motion learning on K400 (without using its action labels), we fine-tune the network on UCF-101 for action recognition, as shown in Table 5. For fair comparisons, we adopt the S3D network as backbone and follow the same experimental setting as used in CoCLR [17]. Interestingly, our approach is capable of learning effective video representation that is comparable with state-of-the-art self-supervised learning methods, even though network pre-training is **not** the main focus of our work. Note that previous work requires pre-extracted optical flow for model training [18, 32] or much larger pre-training datasets such as YouTube8M (YT-8M) [1] to achieve state-of-the-art results.

4.3 Analysis on Motion Features at Different Levels

Here we investigate the different semantics integrated in the motion features learned at different levels. We first conduct a video retrieval experiment using the learned motion features and find that motion features at lower levels tend to retrieve the videos sharing similar elementary movements but lacking of high-level action correlation, as illustrated in Figure 5.

Based on the retrieval results, we can further build a k-nearest-neighbor (KNN) classifier that assigns the query video to the majority action class among its top k nearest neighbors. As shown in Table 6, using low-level motion features can obtain relatively high accuracy for some action classes with apparent moving patterns (e.g., “turning the camera left”). This indicates that the low-level motion features are capable of extracting elementary movements from raw video frames. On the other hand, motion features at higher levels can recognize actions that require finer understanding of high-level motion semantics (e.g., “pulling two ends of smth”). More details are available in the supplementary material.

5 Conclusion

We have presented hierarchical contrastive motion learning, a multi-level self-supervised framework that progressively learns a hierarchy of motion features from raw video frames. A discriminative contrastive loss at each level provides explicit self-supervision for motion learning. This hierarchical design bridges the semantic gap between low-level motion cues and high-level recognition tasks, meanwhile promotes effective fusion of appearance and motion information to finally boost action recognition. Extensive experiments on four benchmarks show that our approach compares favorably against the state-of-the-art methods yet without requiring optical flow or supervised pre-training.

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] Johannes Bill, Hrag Pailian, Samuel J Gershman, and Jan Drugowitsch. Hierarchical structure is employed by humans during visual motion perception. *Journal of Vision*, 2019.
- [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] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, 2005.
- [6] Nieves Crasto, Philippe Weinzaepfel, Karteek Alahari, and Cordelia Schmid. MARS: Motion-augmented RGB stream for action recognition. In *CVPR*, 2019.
- [7] Ali Diba, Vivek Sharma, Luc Van Gool, and Rainer Stiefelhagen. DynamoNet: Dynamic action and motion network. In *ICCV*, 2019.
- [8] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *ICCV*, 2015.
- [9] Lijie Fan, Wenbing Huang, Chuang Gan, Stefano Ermon, Boqing Gong, and Junzhou Huang. End-to-end learning of motion representation for video understanding. In *CVPR*, 2018.
- [10] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. SlowFast networks for video recognition. In *ICCV*, 2019.
- [11] Basura Fernando, Hakan Bilen, Efstratios Gavves, and Stephen Gould. Self-supervised video representation learning with odd-one-out networks. In *CVPR*, 2017.
- [12] Martin A Giese and Tomaso Poggio. Neural mechanisms for the recognition of biological movements. *Nature Reviews Neuroscience*, 2003.
- [13] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Freund, Peter Yianilos, and Moritz Mueller-Freitag. The “something something” video database for learning and evaluating visual common sense. In *ICCV*, 2017.
- [14] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, 2010.
- [15] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006.

- [16] Tengda Han, Weidi Xie, and Andrew Zisserman. Video representation learning by dense predictive coding. In *ICCV Workshop*, 2019.
- [17] Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning. In *ECCV*, 2020.
- [18] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *NeurIPS*, 2020.
- [19] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3D CNNs retrace the history of 2D CNNs and ImageNet? In *CVPR*, 2018.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [21] Asmaa Hosni, Christoph Rhemann, Michael Bleyer, Carsten Rother, and Margrit Gelautz. Fast cost-volume filtering for visual correspondence and beyond. *TPAMI*, 2012.
- [22] Yu Jason, Adam Harley, and Konstantinos Derpanis. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *ECCV*, 2016.
- [23] Boyuan Jiang, Mengmeng Wang, Weihao Gan, Wei Wu, and Junjie Yan. STM: Spatiotemporal and motion encoding for action recognition. In *ICCV*, 2019.
- [24] Dahun Kim, Donghyeon Cho, and Inso Kweon. Self-supervised video representation learning with space-time cubic puzzles. In *AAAI*, 2019.
- [25] Heeseung Kwon, Manjin Kim, Suha Kwak, and Minsu Cho. MotionSqueeze: Neural motion feature learning for video understanding. In *ECCV*, 2020.
- [26] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. In *ICCV*, 2017.
- [27] Ji Lin, Chuang Gan, and Song Han. TSM: Temporal shift module for efficient video understanding. In *ICCV*, 2019.
- [28] Ce Liu. Beyond pixels: Exploring new representations and applications for motion analysis. *PhD Thesis*, 2009.
- [29] William Lotter, Gabriel Kreiman, and David Cox. Deep predictive coding networks for video prediction and unsupervised learning. In *ICLR*, 2017.
- [30] Zhichao Lu, Vivek Rathod, Ronny Votel, and Jonathan Huang. RetinaTrack: Online single stage joint detection and tracking. In *CVPR*, 2020.
- [31] Chenxu Luo, Xiaodong Yang, and Alan Yuille. Self-supervised pillar motion learning for autonomous driving. In *CVPR*, 2021.
- [32] Chenxu Luo, Xiaodong Yang, and Alan Yuille. Exploring simple 3D multi-object tracking for autonomous driving. In *ICCV*, 2021.

- [33] Behrooz Mahasseni, Xiaodong Yang, Pavlo Molchanov, and Jan Kautz. Budget-aware activity detection with a recurrent policy network. In *BMVC*, 2018.
- [34] Yue-Hei Ng, Jonghyun Choi, Jan Neumann, and Larry Davis. ActionFlowNet: Learning motion representation for action recognition. In *WACV*, 2018.
- [35] Simon Niklaus and Feng Liu. Context-aware synthesis for video frame interpolation. In *CVPR*, 2018.
- [36] Aaron Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018.
- [37] AJ Piergiovanni and Michael Ryoo. Representation flow for action recognition. In *CVPR*, 2019.
- [38] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning. *CVPR*, 2021.
- [39] Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xinmei Tian, and Tao Mei. Learning spatiotemporal representation with local and global diffusion. In *CVPR*, 2019.
- [40] Laura Sevilla-Lara, Shengxin Zha, Zhicheng Yan, Vedanuj Goswami, Matt Feiszli, and Lorenzo Torresani. Only time can tell: Discovering temporal data for temporal modeling. In *WACV*, 2021.
- [41] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, 2014.
- [42] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv:1212.0402*, 2012.
- [43] Jonathan Stroud, David Ross, Chen Sun, Jia Deng, and Rahul Sukthankar. D3D: Distilled 3D networks for video action recognition. In *WACV*, 2020.
- [44] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *CVPR*, 2018.
- [45] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Models matter, so does training: An empirical study of CNNs for optical flow estimation. *TPAMI*, 2019.
- [46] Li Tao, Xueting Wang, and Toshihiko Yamasaki. Self-supervised video representation learning using inter-intra contrastive framework. In *ACM Multimedia*, 2020.
- [47] 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.
- [48] Heng Wang, Du Tran, Lorenzo Torresani, and Matt Feiszli. Video modeling with correlation networks. In *CVPR*, 2020.
- [49] Jiangliu Wang, Jianbo Jiao, and Yun-Hui Liu. Self-supervised video representation learning by pace prediction. In *ECCV*, 2021.

- [50] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016.
- [51] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [52] 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.
- [53] Ceyuan Yang, Yinghao Xu, Bo Dai, and Bolei Zhou. Video representation learning with visual tempo consistency. *arXiv:2006.15489*, 2020.
- [54] Xiaodong Yang, Pavlo Molchanov, and Jan Kautz. Multilayer and multimodal fusion of deep neural networks for video classification. In *ACM MM*, 2016.
- [55] Xiaodong Yang, Pavlo Molchanov, and Jan Kautz. Making convolutional networks recurrent for visual sequence learning. In *CVPR*, 2018.
- [56] Xitong Yang, Xiaodong Yang, Ming-Yu Liu, Fanyi Xiao, Larry Davis, and Jan Kautz. STEP: Spatio-temporal progressive learning for video action detection. In *CVPR*, 2019.
- [57] Yue Zhao, Yuanjun Xiong, and Dahua Lin. Recognize actions by disentangling components of dynamics. In *CVPR*, 2018.
- [58] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *ECCV*, 2016.
- [59] Yi Zhu, Zhenzhong Lan, Shawn Newsam, and Alexander Hauptmann. Hidden two-stream convolutional networks for action recognition. In *ACCV*, 2018.
- [60] Mohammadreza Zolfaghari, Kamaljeet Singh, and Thomas Brox. Eco: Efficient convolutional network for online video understanding. In *ECCV*, 2018.