

Beyond Short Clips: End-to-End Video-Level Learning with Collaborative Memories

Xitong Yang^{1*}, Haoqi Fan², Lorenzo Torresani^{2,3}, Larry Davis¹, Heng Wang²

¹University of Maryland, College Park ²Facebook AI ³Dartmouth

{xyang35, lsd}@cs.umd.edu {haoqifan, torresani, hengwang}@fb.com

Abstract

The standard way of training video models entails sampling at each iteration a single clip from a video and optimizing the clip prediction with respect to the video-level label. We argue that a single clip may not have enough temporal coverage to exhibit the label to recognize, since video datasets are often weakly labeled with categorical information but without dense temporal annotations. Furthermore, optimizing the model over brief clips impedes its ability to learn long-term temporal dependencies. To overcome these limitations, we introduce a collaborative memory mechanism that encodes information across multiple sampled clips of a video at each training iteration. This enables the learning of long-range dependencies beyond a single clip. We explore different design choices for the collaborative memory to ease the optimization difficulties. Our proposed framework is end-to-end trainable and significantly improves the accuracy of video classification at a negligible computational overhead. Through extensive experiments, we demonstrate that our framework generalizes to different video architectures and tasks, outperforming the state of the art on both action recognition (e.g., Kinetics-400 & 700, Charades, Something-Something-V1) and action detection (e.g., AVA v2.1 & v2.2).

1. Introduction

In recent years, end-to-end learning of 3D convolutional networks (3D CNNs) has emerged as the prominent paradigm for video classification [2, 5, 7, 9, 10, 22, 33, 39, 41–43, 45, 47, 53]. Steady improvements in accuracy have come with the introduction of increasingly deeper and larger networks. However, due to their high computational cost and large memory requirements, most video models are optimized at each iteration over short, fixed-length clips rather than the entire video.

Although widely used in modern video models, the clip-level learning framework is sub-optimal for video-level

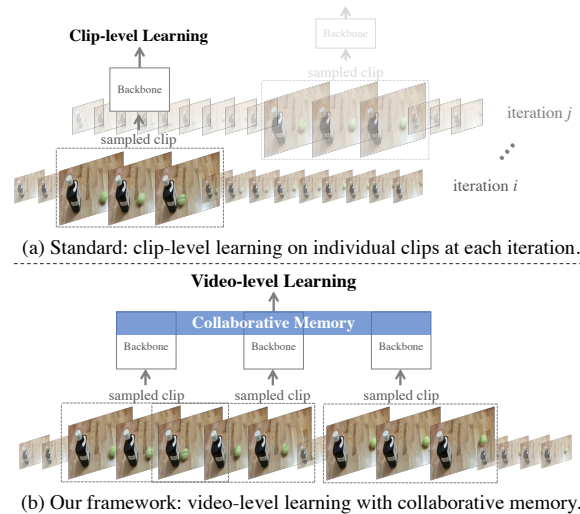


Figure 1: Clip-level learning vs. our proposed end-to-end video-level learning framework. (Action label: *something being deflected from something*.)

classification. First, capturing long-range temporal structure beyond short clips is not possible as the models are only exposed to individual clips during training. Second, the video-level label may not be well represented in a brief clip, which may be an uninformative segment of the video or include only a portion of the action, as shown in Figure 1(a). Thus, optimizing a model over individual clips using video-level labels is akin to training with noisy labels. Recent attempts to overcome these limitations include methods that build a separate network on top of the clip-based backbone [20, 50, 56]. However, these approaches either cannot be trained end-to-end with the backbone (*i.e.*, the video model is optimized over pre-extracted clip-level features) or require ad-hoc backbones which hinder their application in the current landscape of evolving architectures.

In this paper, we propose an end-to-end learning framework that optimizes the classification model using video-level information collected from multiple temporal locations of the video, shown in Figure 1(b). Our approach

*Work done during an internship at Facebook AI.

hinges on a collaborative memory mechanism that accumulates video-level contextual information from multiple clips sampled from the video. Within the same training iteration, this contextual information is shared back with all the clips to enhance the individual clip representations. The collaborative memory allows the model to capture long-range temporal dependencies beyond individual short clips by generating clip-specific memories that encode the relation between each local clip and the global video-level context.

Our experiments demonstrate that the proposed training framework is effective and generic. Specifically, our approach does not make any assumption about the backbone architecture. We empirically show that it consistently yields significant gains in accuracy when applied to different state-of-the-art architectures (e.g. SlowFast [10], R(2+1)D [43], I3D-NL [47]). We also introduce and compare several design variants of the collaborative memory. Furthermore, we demonstrate that the accuracy improvements come at a negligible computational overhead and without an increase in memory requirements. Finally, we show that our framework can be extended to action detection where it yields significant improvements without requiring extra information, such as optical flow and object detection predictions, which are commonly used in previous work [38, 40]. We summarize our major contributions as follows:

- A new framework that enables end-to-end learning of video-level dependencies for clip-based models.
- A new collaborative memory mechanism that facilitates information exchange across multiple clips. We explore different design choices and provide insights about the optimization difficulties.
- Experiments demonstrating that our collaborative memory framework generalizes to different backbones and tasks, producing state of the art results for action recognition and detection.

2. Related Work

Clip-Level Video Architectures. Since the introduction of 3D CNNs [2, 22, 41] to video classification, new architectures [5, 7, 9, 33, 39, 43, 45, 47, 53] have been proposed to learn better spatiotemporal representations. Besides models aimed at improving accuracy, several architectures have been proposed to achieve better performance/cost trade-offs [9, 24, 27, 32, 42, 51, 57]. Another line of research involves the design of multiple-stream networks [5, 10, 11, 36, 46, 53, 54], where each stream consumes a different type of input, e.g., RGB or optical flow. Besides CNNs, transformer-based models, e.g., TimeSformer [3], also show promising results. Unlike prior work focused on the design of clip-level architecture, our paper proposes a new framework to learn long-range dependencies using existing clip-level models. As we do not make any assumption

about the clip-level architecture, our framework generalizes to different backbones and enables end-to-end training of clip models with video-level contextual information.

Video-Level Classification. Several attempts have been made to overcome the limitations of the single-clip training framework. Timeception [20] uses multi-scale temporal convolutions to cover different temporal extents for long-range temporal modeling. Timeception layers are trained on top of a frozen backbone. Both TSN [46] and ECO [57] divide the input video into segments of equal size and randomly sample a short snippet or a single frame from each segment to provide better temporal coverage during training. As the GPU memory cost grows linearly w.r.t. the number of segments, TSN and ECO adopt lightweight 3D CNNs or even 2D CNNs as backbones in order to process multiple segments simultaneously. These simple backbones limit the performance of the framework. In addition, TSN uses averaging to aggregate the predictions from different segments, whereas we propose a dedicated memory mechanism to model the video-level context. FASTER [56] and SCSampler [25] explore strategies to limit the detrimental impact of applying video-level labels to clips and to save computational cost.

Another related work to our approach is LFB [50]. It leverages context features from other clips to augment the prediction on the current clip. Unlike our approach, context features stored in LFB are *pre-computed* using a *separate* model. As a result, the context features cannot be updated during the training and the model used to extract these context features is not optimized for the task. In contrast, our framework is end-to-end trainable and the accumulated contextual information can back-propagate into the backbone architecture. Note that storing context features is infeasible for large-scale video datasets, e.g., Kinetics, and LFB is mainly designed for action detection applications.

Learning With Memories. Memory mechanisms [1, 15, 37, 49] have been widely used in Recurrent Neural Networks for language modeling in order to learn long-term dependencies from sequential text data. Specifically, memory networks [49] have been proposed for question answering (QA), while Sukhbaatar *et al.* [37] have introduced a strategy enabling end-to-end learning of these models. RWMN [31] has extended the QA application on movie videos. Grave *et al.* [15] have proposed to store past hidden activations as a memory that can be accessed through a dot product with the current hidden activation.

These works are similar in spirit to our approach, but our application is in a different domain with different constraints and challenges. Moreover, our collaborative memory mechanism is designed to capture the interactions among the samples, is extremely lightweight and memory-friendly, and is suitable for training computationally intensive video models.

3. End-to-End Video-Level Learning with Collaborative Memory

We start with an overview of the proposed framework, then present different designs for the collaborative memory. We conclude with a discussion of implementation strategies to cope with the GPU memory constraint.

3.1. Overview of the Proposed Framework

Given a video recognition architecture (e.g., I3D [5], R(2+1)D [43], SlowFast [10]) that operates on short, fixed-length clips, our goal is to perform video-level learning in an end-to-end manner. In particular, we aim to optimize the *clip-based* model using *video-level* information collected from the whole video. To achieve this, we start by sampling multiple clips from the video within the same training iteration in order to cover different temporal locations of the video. The clip-based representations generated from multiple clips are then accumulated via a collaborative memory mechanism that captures interactions among the clips and builds video-level contextual information. After that, clip-specific memories are generated to enhance the individual clip-based representations by infusing the video-level information into the backbone. Finally, the sampled clips are jointly optimized with a video-level supervision applied to the consensus of predictions from multiple clips.

Multi-clip sampling. Given a video $\mathcal{V} = \{I_0, \dots, I_{T-1}\}$ with T frames, we sample N clips $\{\mathcal{C}_0, \dots, \mathcal{C}_{N-1}\}$ from the video at each training iteration. Each short clip $\mathcal{C}_n = \{I_{t_n}, \dots, I_{t_n+L-1}\}$ consists of L consecutive frames randomly sampled from the full-length video where t_n indicates the index of the start frame. N is a hyper-parameter that can be decided based on the ratio between the video length and the clip length to ensure sufficient temporal coverage. The sampled clips are then fed to the backbone to generate clip-based representations $\{X_n\}_{n=0}^{N-1}$, where $X_n = f(\mathcal{C}_n)$, and f represents the clip-level backbone. In the traditional clip-level classification, X_n is directly used to perform the final prediction via a classifier h : $\mathbf{y}_n = h(X_n) = h(f(\mathcal{C}_n))$, where \mathbf{y}_n is the vector of classification scores.

Collaborative memory. Our approach hinges on a collaborative memory mechanism that accumulates information from multiple clips for learning video-level dependencies and then shares this video-level context back with the individual clips, as illustrated in Figure 2. Specifically, the collaborative memory involves two stages:

- **Memory interactions:** A global memory of the whole video is constructed by accumulating information from multiple clips: $\mathcal{M} = Push(\{X_n\}_{n=0}^{N-1})$. The global memory is then shared back with the individual clips in order to generate clip-specific memories: $M_n = Pop(\mathcal{M}, X_n)$.

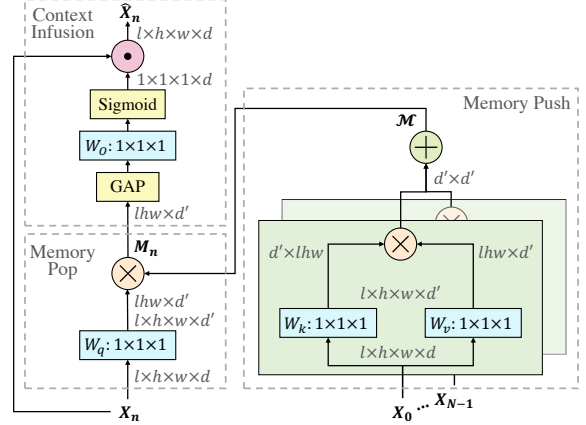


Figure 2: Collaborative memory with associative memory and feature gating. The feature maps are shown as the shape of their tensors, e.g., $l \times h \times w \times d$. GAP denotes global average pooling. \otimes and \odot indicate matrix and elementwise multiplication, respectively.

- **Context infusion:** The individual clip-based representations are infused with video-level context. This is done by means of a gating function g that enhances each clip representation with the information stored in the clip-specific memory: $\hat{X}_n = g(X_n, M_n)$.

Video-level supervision. To facilitate the joint optimization over multiple clips, we apply a video-level loss that takes into account the clip-level predictions as well as the video-level prediction aggregated from all N sampled clips. Formally, we first aggregate the clip-level predictions via average pooling over N clips: $H = \frac{1}{N} \sum_{n=0}^{N-1} h(\hat{X}_n) = \frac{1}{N} \sum_{n=0}^{N-1} h(g(f(\mathcal{C}_n), M_n))$. Then our video-level loss can be written as

$$\mathcal{L}_{video} = \frac{1}{N} \sum_{n=0}^{N-1} \mathcal{L}(h(\hat{X}_n)) + \alpha \mathcal{L}(H). \quad (1)$$

\mathcal{L} denotes the cross-entropy loss for classification and α is the weight to balance the two terms which account for the clip-level losses and the video-level loss aggregated from all N clips. All the parameters (i.e., f , g , and h) are optimized end-to-end w.r.t. this objective.

3.2. Collaborative memory

Our idea of collaborative memory is generic and can be implemented in a variety of ways. In this section we introduce a few possible designs. We empirically evaluate these different options in Section 4.3.

Memory interactions. The design of memory interactions should follow two principles: 1) The memory footprint for storing the global memory should be manageable; 2) Interactions with the memory should be computationally efficient. The first principle implies that the memory

consumption should not grow with the number of clips N . Thus, simply storing all clip-based features is not feasible.

– **Average Pooling:** Let the clip-based representation X_n be a $k \times d$ matrix, where k is the spatial-temporal resolution (i.e., $height \times width \times length$) and d is the number of channels. A simple strategy is to perform a global average pooling over all sampled clips: $\mathcal{M} = Push(\{X_n\}_{n=0}^{N-1}) = Pool(\{X_n W_I\}_{n=0}^{N-1})$. $W_I \in \mathbb{R}^{d \times d'}$ is a learnable weight matrix to reduce the dimensionality from d to d' . This global information can be simply shared back with all the clips: $M_n = Pop(\mathcal{M}, X_n) = \mathcal{M}$.

– **Associative Memory:** Although avg/max pooling is capable of collecting information from multiple clips, it fails to capture the inter-clip dependencies and the clip-specific information cannot be retrieved from the global memory \mathcal{M} . This motivates us to design a new mechanism that enables the retrieval of clip-specific information when needed. Inspired by associative networks [1, 18], we propose to accumulate the clip-level features using the outer product operator to generate the global memory \mathcal{M} :

$$Push(\{X_n\}_{n=0}^{N-1}) = \frac{1}{N} \sum_{n=0}^{N-1} (X_n W_k)^T (X_n W_v). \quad (2)$$

Given the n -th clip, we obtain its clip-specific memory by:

$$M_n = Pop(\mathcal{M}, X_n) = (X_n W_q) \mathcal{M}, \quad (3)$$

where $W_k, W_v, W_q \in \mathbb{R}^{d \times d'}$ are learnable weight matrices for memory interactions and dimension reduction. Note that this memory design can be viewed as implementing a form of video-level inter-clip attention, where the clip-based representation X_n attends to features generated from all sampled clips of the video in proportion to their similarities: $M_n = \frac{1}{N} \sum_{m=0}^{N-1} [(X_n W_q)(X_m W_k)^T] (X_m W_v)$. However, unlike the self-attention mechanism [44, 47], our design is more efficient in both computation and memory consumption as it does not require to store all clip-level features or perform pairwise comparison between all the clips.

Context Infusion. One way to incorporate the clip-specific memory M_n with the clip-level features X_n is through a residual connection: $\hat{X}_n = M_n W_O + X_n$, where $W_O \in \mathbb{R}^{d' \times d}$ is a linear transformation to match the feature dimensionality. However, as we will show in our experiments (Figure 5), this design tends to overfit to the clip-specific memory during training and leads to inferior performance. As M_n stores much more information than a single clip-level feature X_n , the model mostly relies on M_n during training and makes little use of X_n .

In light of the above observation, we propose to infuse context information into the clip-level features through a *feature gating* operation. Rather than allowing the model to directly access the clip-specific memory M_n , feature gating forces the model to recalibrate the strengths of different

clip-level features using the contextual information. Formally, the enhanced features are computed as

$$\hat{X}_n = \left(J + \sigma(\hat{M}_n W_O) \right) \odot X_n, \quad (4)$$

where σ is the sigmoid activation function, \odot is the element-wise multiplication and J is an all-ones matrix for residual connection. \hat{M}_n is obtained by averaging the spatial and temporal dimensions of M_n : $\hat{M}_n = GAP(M_n)$. Our feature gating operation can be considered as a channel-wise attention mechanism similar to context gating [30, 53] and the SE block [19]. However, the attention weights in our method are generated by video-level contextual information, instead of self-gating values that capture channel-wise relationships within the same clip. Experimental results show that our proposed feature gating design alleviates the optimization difficulties during training and enables a more effective use of the video-level contextual information.

3.3. Coping with the GPU Memory Constraint

A challenge posed by video-level learning is the need to jointly optimize over multiple clips under a fixed and tight GPU memory budget. In this section, we discuss two simple implementations that allow end-to-end training of video-level dependencies under this constraint.

Batch reduction. Let B be the size of the mini-batch of videos used for traditional clip-level training. Our approach can be implemented under the same GPU memory budget by reducing the batch size by a factor of N : $\hat{B} = round(B/N)$. This allows us to load into the memory N clips for each of the B/N different videos. In order to improve the clip diversity for updating the batch-norm [21] parameters within a mini-batch, we propose to calculate batch-norm statistics using only the clips from different videos. Although this implementation cannot handle arbitrarily large N , it is simple, efficient and we found it applicable to most settings in practice. For example, a typical choice of batch size for training clip-based models is a 8-GPU machine with $B = 64$; our approach can be implemented under this memory setup by using in each iteration $\hat{B} = 16$ different videos, and by sampling from each of them $N = 4$ clips.

Multi-iteration. Instead of directly loading N clips into one mini-batch, we can also unroll the training of a video into N iterations. Each iteration uses one of the N clips. This implementation is memory-friendly and consumes the same amount of memory as the standard single-clip training framework. It allows us to process arbitrarily long videos with arbitrarily large N . When incorporating the collaborative memory, we simply perform a two-scan process: the first scan generates the clip-based features to build the global memory \mathcal{M} and the second scan generates the classification output of each clip conditioned on \mathcal{M} .





Figure 3: **Video-level accuracy** on Kinetics-400 with CM. The horizontal axis shows the number of clips (N) used for training CM. Note that all models use 30 crops for inference. Figure 4: **Clip-level accuracy** on Kinetics-400 at 10 different temporal locations within the video. N indicates the number of clips used for training CM.

Model	Baseline	Ours	Δ	FLOPs
Slow-only-50 8×8 [10]	74.4	76.8	+2.4	1.03 \times
I3D-50+NL 32×2 [47]	74.9	77.5	+2.4	1.02 \times
R(2+1)D-50 16×2 [43]	75.7	78.0	+2.3	1.01 \times
SlowFast-50 4×16 [10]	75.6	77.8	+2.2	1.02 \times
SlowFast-50 8×8 [10]	76.8	78.9	+2.1	1.03 \times

Table 1: Generalization to **different backbone architectures**. We report the video-level accuracy on Kinetics-400 for both standard clip-level training (“Baseline”) and video-level training with CM (“Ours”).

4. Experiments

To demonstrate the advantages of our end-to-end video-level learning framework, we conduct extensive experiments on four action recognition benchmarks with different backbone architectures. We implement our models and conduct the experiments using the PySlowFast codebase [8].

4.1. Experimental Setup

Datasets. Kinetics [23] (K400 & K700) is one of the most popular datasets for large-scale video classification. Charades [35] is a multi-label dataset with long-range activities. Something-Something-V1 [14] is a dataset requiring good use of temporal relationships for accurate recognition. Following the standard protocol, we use the training set for training and report top-1 accuracy on the validation set.

Backbones. We evaluate our framework using different backbone architectures including I3D [47], R(2+1)D [43, 45], Slow-only [10] and SlowFast [10], optionally augmented with non-local blocks (NL) [47]. We attach the proposed collaborative memory to the last convolutional layer of these backbones for joint training.

Training. We first train the backbones by themselves following their original schedules [10, 45], then re-train the backbones in conjunction with our collaborative memory for video-level learning. When training on Kinetics, we use synchronous SGD with a cosine learning schedule [28]. For Charades and Something-Something-V1, we follow the recipe from PySlowFast [8] and initialize the network weights from the models pre-trained on Kinetics, since these two datasets are relatively small. For the video-level training, we employ the batch reduction strategy to handle the GPU memory constraint by default and apply the linear scaling rule [13] to adjust the training schedule accordingly.

Inference. Following [10, 47], we uniformly samples 3×10 crops from each video for testing (*i.e.*, 3 spatial crops and 10 temporal crops). The global memory \mathcal{M} is aggregated from 10 spatially centered crops and shared for the

inference of all 30 crops. We employ the multi-iteration approach from Section 3.3 during inference to overcome the GPU memory constraint. The softmax scores of all 30 clips are averaged for the final video-level prediction.

4.2. Evaluating Collaborative Memory

For all the experiments in this section we use the associative version of the collaborative memory with feature gating since, as demonstrated in ablation studies (Section 4.3), this design provides the best results.

Effectiveness of video-level learning. We begin by presenting an experimental comparison between our proposed video-level learning and the standard clip-level training applied to the same architecture. For this evaluation we use the Slow-only backbone with 50 layers, which can be considered as a 3D ResNet [17]. In order to investigate the impact of temporal coverage on video-level learning, we train models using different numbers of sampled clips per video: $N \in \{1, 3, 5, 7, 9\}$. $N = 1$ corresponds to the conventional clip-level training, as we only sample one clip per video. In such case the collaborative memory (CM) is limited to perform “self-attention” within the single clip. For $N > 1$, CM captures video-level information across the N clips.

Figure 3 shows the *video*-level accuracy achieved by the different models, all using the same testing setup of 3×10 crops per video. Note that under this setting all models “look” at the same number of clips for each video in testing. As shown in Figure 3, our CM framework significantly improves the video-level accuracy. For example, when the clip length is 8×8 (8 frames with a temporal stride of 8), using CM with $N = 9$ yields a remarkable 2.6% improvement compared with training using a single clip (74.5% vs. 77.1%). When the clip has a shorter length (*i.e.*, 8×2), our method gives an even larger gain, 3.2% (73.2% vs. 76.4%). As expected, the improvement from our method increases with the number of sampled clips N . The performance saturates when $N \geq 7$. To keep the training time more manageable, we use $N = 5$ by default.

Multi-clip	Memory	End-to-end	Top-1
	✓	✓	74.5
✓		✓	75.5
✓	✓		75.9
✓	✓	✓	76.8

(a) Evaluating **different components** of our video-level learning framework.

Model	Stage-wise	Top-1
Slow-only		76.1
	✓	76.8
R(2+1)D		77.7
	✓	78.0

(d) **Stage-wise training** vs. training everything from scratch.

Setting	Associative	Gating	Top-1
Multi-clip (w/o memory)			75.5
CM (avgpool)		✓	75.8
CM (residual)	✓		76.0
CM (default)	✓	✓	76.8

(b) Comparing **different designs** of our collaborative memory mechanism.

Model	Batch reduction	Multi-iteration	Top-1
Slow-only		✓	76.6
	✓		76.8
R(2+1)D		✓	77.9
	✓		78.0

(e) Comparing different ways of training CM: **batch reduction** vs. multi-iteration.

Setting	#Param.	Top-1
$\alpha = 1$	49.2 M	76.8
$\alpha = 2$	40.9 M	76.8
$\alpha = 4$	36.7 M	76.8
$\alpha = 8$	34.6 M	76.4

(c) Varying **channel reduction ratio** $\alpha = d/d'$.

Model	Temporal stride				CM
	2	4	8	16	
Slow-only	73.2	74.3	74.4	74.4	76.8
R(2+1)D	75.7	76.4	75.0	72.2	78.0

(f) Comparing CM with backbones using clips with **large temporal strides**.

Table 2: Ablation experiments on Kinetics-400. Top-1 video-level accuracy (%) is reported. Unless otherwise stated, we use Slow-only [10] with 50 layers and the input clip length is 8×8 . R(2+1)D is also 50 layers with a clip length of 16×2 .

Figure 4 shows the clip-level accuracy at different temporal locations of a video, according to their temporal order. When $N = 1$, the clips from the middle of the video have much higher accuracy than the clips from the beginning or the end of the video, as the middle clips tend to include more relevant information. CM significantly improves the clip-level accuracy by augmenting each clip with video-level context information (*i.e.*, $N \geq 3$), especially for clips near the boundary of the video. This is a clear indication that our memory mechanism is capable of capturing video-level dependencies and sharing them effectively with the clips within the video to boost the recognition accuracy.

Generalization to different backbones. As we do not make any assumption about the backbone, our video-level end-to-end learning framework can be seamlessly integrated with different architectures and input configurations (*e.g.*, clip length, sampling stride, *etc.*). As shown in Table 1, CM produces consistent video-level accuracy gains of over 2% on top of state-of-the-art video models, including I3D with non-local blocks [47], the improved R(2+1)D network [43, 45] and the recent SlowFast network [10]. Note that we achieve these improvements with only negligible additional inference cost, about 1-3% more FLOPs compared to the backbone themselves.

4.3. Ablation Studies

Assessing the components in our framework. Unlike most prior work on video-level modeling [20, 50], our framework is end-to-end trainable. To show the benefits of end-to-end learning in improving the backbone, we conduct an ablation that freezes the parameters of the backbone and only updates the parameters from the collaborative mem-

ory and the FC layers for classification. As shown in Table 2a, end-to-end learning improves the performance by 1.1% compared with learning video-level aggregation on top of the frozen backbones (76.8% vs. 75.9%).

Table 2a also shows the result of video-level learning without using CM. This is done by optimizing multiple clips per video but without sharing any information across the clips. Interestingly, this simple setup also delivers a good improvement over the single-clip learning baseline (75.5% vs. 74.5%). The gain comes from the joint optimization over multiple clips of a video, which again confirms the importance of video-level learning for classification. Our CM framework achieves the best performance with all the components enabled.

Collaborative memory design. Our default design uses the associative memory for memory interactions and a feature gating operation for context infusion. In Table 2b, we explore other design choices by replacing the associative memory with average pooling or substituting the feature gating with a residual connection.

We observe that all these variants offer improvement over the naïve video-level learning setup without the memory. However, the gain provided by average pooling is relatively small, which is not surprising given that there is no the inter-clip interaction. While we also witness a performance drop by removing the feature gating operation, the reason behind it is different. As shown by the training/validation error curves in Figure 5, the model without feature gating achieves lower training error but higher validation error. This suggests that the model degenerates during training due to over-fitting to the video-level context.

In Table 2c we ablate the number of channels used in




Figure 5: Video-level training/validation errors on Kinetics-400 for different designs of the collaborative memory.

CM (d' in Eq. 2, 3), which can be controlled by the channel reduction ratio $\alpha = d/d'$. We can see that the results remain unchanged as long as the reduction ratio is reasonable ($\alpha \leq 4$). We use $\alpha = 4$ as the default value in our experiments since it introduces fewer extra parameters.

Training strategies. Recall that we adopt a stage-wise training strategy: the backbones are first trained using standard clip-level training recipes and then re-trained in conjunction with CM for video-level learning. In Table 2d, we compare this strategy with training everything from scratch (equivalent training iterations are used for both strategies). Experiments on two different backbones show that training everything from scratch yields slightly worse results. We hypothesize that stage-wise training allows the optimization in the second stage to focus on effective long-range modeling thanks to the well-initialized backbone. We note that non-local networks are also trained in a stage-wise way.

We also compare the two methods to cope with the GPU memory constraint (Section 3.3). As shown in Table 2e, the accuracy of the two methods is almost the same and the difference is within the margin of randomness, which makes sense as the two methods are technically identical.

Limitations of temporal striding. One simple way to increase the temporal coverage of a video model is to use larger temporal strides when sampling the frames of the input clips. We compare our video-level learning framework with this strategy in Table 2f. Note that we keep the temporal strides of CM the same as the original backbones, *i.e.*, 8 frames for Slow-only and 2 frames for R(2+1)D.

We can see that increasing the temporal coverage through striding yields only a modest gain in accuracy. Notably, using a very large stride even hurts performance for some architectures like R(2+1)D. In contrast, our approach can learn long-range dependencies and improves the performances of short clip-based backbones by large margins.

4.4. Comparison with the State of the Art

Previous experimental results are from Kinetics-400. To demonstrate that our method can generalize to different datasets, we further evaluate our method on Kinetics-700 [4], Charades [35] and Something-Something-V1 [14].

Methods	Pretrain	Only RGB	GFLOPs × crops	Top-1
I3D [5]	ImageNet	✗	216×N/A	75.7
S3D-G [53]	ImageNet	✗	142.8×N/A	77.2
LGD-3D-101 [34]	ImageNet	✗	N/A	81.2
I3D-101+NL [47]	ImageNet	✓	359×30	77.7
ip-CSN-152 [42]	Sports1M	✓	109×30	79.2
CorrNet-101	Sports1M	✓	224×30	81.0
MARS+RGB [6]	none	✓	N/A	74.8
DynamoNet [7]	none	✓	N/A	77.9
CorrNet-101 [45]	none	✓	224×30	79.2
SlowFast-101 8×8 [10]	none	✓	106×30	77.9
SlowFast-101 16×8 [10]	none	✓	213×30	78.9
SlowFast-101+NL 16×8 [10]	none	✓	234×30	79.8
Ours (R(2+1)D-101 32×2)	none	✓	243×30	80.5
Ours (SlowFast-101 8×8)	none	✓	128×30	80.0
Ours (SlowFast-101+NL 8×8)	none	✓	137×30	81.4

Table 3: Comparison with the state-of-the-art on Kinetics-400.

Methods	Pretrain	GFLOPs × crops	Top-1
SlowFast-101+NL 8×8 [10]	K600	115×30	70.6
SlowFast-101+NL 16×8 [10]	K600	234×30	71.0
SlowFast-50 4×16*	K600	36×30	66.1
SlowFast-101 8×8*	K600	126×30	69.2
SlowFast-101+NL 8×8*	K600	135×30	70.2
Ours (SlowFast-50 4×16)	K600	37×30	68.3
Ours (SlowFast-101 8×8)	K600	128×30	70.9
Ours (SlowFast-101+NL 8×8)	K600	137×30	72.4

Table 4: Comparison with the state-of-the-art on Kinetics-700. * indicates results reproduced by us.

Among them, Charades has longer-range activities (30 seconds on average), whereas Something-Something-V1 includes mostly human-object interactions. We compare the results with the state of the art in Table 3, 4, 5 and 6. Our proposed CM framework yields consistent gains of over 2% for different variants of SlowFast on all datasets. These improvements are very significant given that SlowFast is among the best video backbones.

On Kinetics-400 and Kinetics-700, our method establishes a new state of the art, as shown in Table 3 and 4. Notably, we achieve these results without pre-training on other datasets or using optical flow. Similarly, our method outperforms the state of the art on both Charades (in Table 5) and Something-Something-V1 (in Table 6). On Charades, our CM framework yields more than 3% gains (*e.g.*, 44.6% vs 41.3%). This demonstrates that CM performs even better on datasets that have longer videos and require longer-term temporal learning. Note that our method significantly out-

Methods	Pretrain	GFLOPs × crops	Top-1
TRN [55]	ImageNet	N/A	25.2
I3D-101+NL [47]	ImageNet+K400	544 × 30	37.5
STRG [48]	ImageNet+K400	630 × 30	39.7
Timeception [20]	K400	N/A	41.1
LFB (I3D-101+NL) [50]	K400	N/A	42.5
SlowFast-101+NL [10]	K400	234×30	42.5
AVSlowFast-101+NL [52]	K400	278×30	43.7
SlowFast-50 16×8*	K400	131×30	39.4
SlowFast-101+NL 16×8*	K400	273×30	41.3
Ours (SlowFast-50 16×8)	K400	135×30	42.9
Ours (SlowFast-101+NL 16×8)	K400	277×30	44.6

Table 5: Comparison with the state-of-the-art on Charades. * indicates results reproduced by us.

Methods	Pretrain	Only RGB	Top-1
S3D-G [53]	ImageNet	✗	48.2
ECO [57]	none	✗	49.5
Two-stream TSM [27]	ImageNet	✗	52.6
MARS+RGB+Flow [6]	K400	✗	53.0
NL I3D-50+GCN [48]	ImageNet	✓	46.1
GST-50 [29]	ImageNet	✓	48.6
MSNet [26]	ImageNet	✓	52.1
CorrNet-101 [45]	Sports1M	✓	53.3
SlowFast-50 8×8*	K400	✓	50.1
SlowFast-101+NL 8×8*	K400	✓	51.2
Ours (SlowFast-50 8×8)	K400	✓	52.3
Ours (SlowFast-101+NL 8×8)	K400	✓	53.7

Table 6: Comparison with the state-of-the-art on Something-Something-V1. * indicates results reproduced by us.

performs other recent work on long-range temporal modeling (e.g., Timeception [20], LFB [50]).

4.5. Collaborative Memory for Action Detection

In this section, we show that the benefits of our framework also extend to the task of action detection. We evaluate our method on AVA [16], which includes 211k training and 57k validation video segments. AVA v2.2 provides more consistent annotations than v2.1 on the same data. We report mean average precision (mAP) over 60 frequent classes on the validation set following the standard protocol.

Adaptation to action detection. Adapting our approach to action detection is straightforward. Instead of randomly sampling multiple clips from the whole video, we sample clips within a certain temporal window $t_n \in [t-w, t+w]$ to detect action at time t , where t_n indicates the center frame of the n th sampled clip and $2w + 1$ is the window size. As AVA includes sparse annotations at one frame per sec-

Methods	Pretrain mAP	Methods	Pretrain mAP
ACRN [38]	K400 17.4 [†]	AVSF-101 8×8 [52]	K400 28.6 [†]
AVSF-50 4×16 [52]	K400 27.8 [†]	AIA(SF-50 4×16) [40]	K700 29.8 [†]
AT (I3D) [12]	K400 25.0	AIA(SF-101 8×8) [40]	K700 32.3 [†]
LFB(R50+NL) [50]	K400 25.8	SF-101+NL 8×8 [10]	K600 29.0
R50+NL* [50]	K400 23.6	SF-50 4×16* [10]	K700 26.9
SF-50 4×16* [10]	K400 23.6	SF-101 8×8* [10]	K700 29.0
Ours (R50+NL)	K400 26.3	Ours (SF-50 4×16)	K700 29.8
Ours (SF-50 4×16)	K400 25.8	Ours (SF-101 8×8)	K700 31.6

(a)

(b)

Table 7: Comparison with SOTA on (a) AVA v2.1 and (b) v2.2. [†] indicates results with extra information other than RGB frames, such as optical flow, audio and objection detection predictions. * indicates results reproduced by us.

ond, we simply sample the clips with a one-second stride such that the sampled clips are centered at frames with annotations. In this way, the temporal window size increases accordingly as we use a larger number of clips N during training. After that, we jointly optimize these sampled clips with their own annotations. The collaborative memory is used to share long-range context information among sampled clips. We use $N = 9$ in our experiments and follow the schedule in AIA [40] for model training.

Quantitative results. We compare CM with the state of the art on AVA in Table 7. Although the CM framework is not specifically designed for action detection, it achieves results comparable with the state of the art. In particular, CM yields a consistent improvement of more than 2% for different backbone networks (e.g., 2.9% gain for SlowFast-50 4×16 backbone on AVA v2.2). This demonstrate that we can effectively extend our method to the detection task and achieve significant improvements as well. Note that our method also outperforms LFB [50] when using the same backbone (i.e., R50-I3D+NL) on AVA v2.1 (26.3% vs. 25.8%).

5. Conclusions

We have presented an end-to-end learning framework that optimizes classification models using video-level information. Our approach hinges on a novel collaborative memory mechanism that accumulates contextual information from multiple clips sampled from the video and then shares back this video-level context to enhance the individual clip representations. Long-range temporal dependencies beyond short clips are captured through the interactions between the local clips and the global memory. Extensive experiments on both action recognition and detection benchmarks show that our framework significantly improves the accuracy of video models at a negligible computational overhead.

References

- [1] Jimmy Ba, Geoffrey E Hinton, Volodymyr Mnih, Joel Z Leibo, and Catalin Ionescu. Using fast weights to attend to the recent past. In *NeurIPS*, 2016. 2, 4
- [2] Moez Baccouche, Franck Mamalet, Christian Wolf, Christophe Garcia, and Atilla Baskurt. Sequential deep learning for human action recognition. In *International Workshop on Human Behavior Understanding*, pages 29–39. Springer, 2011. 1, 2
- [3] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? *arXiv preprint arXiv:2102.05095*, 2021. 2
- [4] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019. 7
- [5] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 1, 2, 3, 7
- [6] Nieves Crasto, Philippe Weinzaepfel, Karteek Alahari, and Cordelia Schmid. Mars: Motion-augmented rgb stream for action recognition. In *CVPR*, 2019. 7, 8
- [7] Ali Diba, Vivek Sharma, Luc Van Gool, and Rainer Stiefelhagen. Dynamonet: Dynamic action and motion network. *ICCV*, 2019. 1, 2, 7
- [8] Haoqi Fan, Yanghao Li, Bo Xiong, Wan-Yen Lo, and Christoph Feichtenhofer. Pyslowfast. <https://github.com/facebookresearch/slowfast>, 2020. 5
- [9] Christoph Feichtenhofer. X3D: Expanding architectures for efficient video recognition. In *CVPR*, 2020. 1, 2
- [10] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019. 1, 2, 3, 5, 6, 7, 8
- [11] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, 2016. 2
- [12] Rohit Girdhar, Joao Carreira, Carl Doersch, and Andrew Zisserman. Video action transformer network. In *CVPR*, 2019. 8
- [13] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. 5
- [14] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Freund, Peter Yianilos, Moritz Mueller-Freitag, et al. The something something video database for learning and evaluating visual common sense. In *ICCV*, 2017. 5, 7
- [15] Edouard Grave, Armand Joulin, and Nicolas Usunier. Improving neural language models with a continuous cache. *arXiv preprint arXiv:1612.04426*, 2016. 2
- [16] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *CVPR*, 2018. 8
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [18] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 1982. 4
- [19] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018. 4
- [20] Noureldien Hussein, Efstratios Gavves, and Arnold WM Smeulders. Timeception for complex action recognition. In *CVPR*, 2019. 1, 2, 6, 8
- [21] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 4
- [22] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3D convolutional neural networks for human action recognition. *TPAMI*, 2013. 1, 2
- [23] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 5
- [24] Okan Köpüklü, Neslihan Kose, Ahmet Gunduz, and Gerhard Rigoll. Resource efficient 3d convolutional neural networks. In *ICCV Workshop*, 2019. 2
- [25] Bruno Korbar, Du Tran, and Lorenzo Torresani. SCSampler: Sampling salient clips from video for efficient action recognition. In *ICCV*, 2019. 2
- [26] Heeseung Kwon, Manjin Kim, Suha Kwak, and Minsu Cho. Motionsqueeze: Neural motion feature learning for video understanding. In *ECCV*, 2020. 8
- [27] Ji Lin, Chuang Gan, and Song Han. TSM: Temporal shift module for efficient video understanding. In *ICCV*, 2019. 2, 8
- [28] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *ICLR*, 2017. 5
- [29] Chenxu Luo and Alan L Yuille. Grouped spatial-temporal aggregation for efficient action recognition. In *ICCV*, 2019. 8
- [30] Antoine Miech, Ivan Laptev, and Josef Sivic. Learnable pooling with context gating for video classification. *arXiv preprint arXiv:1706.06905*, 2017. 4
- [31] Seil Na, Sangho Lee, Jisung Kim, and Gunhee Kim. A read-write memory network for movie story understanding. In *ICCV*, 2017. 2
- [32] A. J. Piergiovanni, Anelia Angelova, and Michael S. Ryoo. Tiny video networks. *CoRR*, abs/1910.06961, 2019. 2
- [33] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3d residual networks. In *ICCV*, 2017. 1, 2
- [34] Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xinmei Tian, and Tao Mei. Learning spatio-temporal representation with local and global diffusion. In *CVPR*, 2019. 7
- [35] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, 2016. 5, 7

- [36] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, 2014. 2
- [37] Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. End-to-end memory networks. In *NeurIPS*, 2015. 2
- [38] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, and Cordelia Schmid. Actor-Centric relation network. In *ECCV*, 2018. 2, 8
- [39] Lin Sun, Kui Jia, Dit-Yan Yeung, and Bertram E. Shi. Human action recognition using factorized spatio-temporal convolutional networks. In *ICCV*, 2015. 1, 2
- [40] Jiajun Tang, Jin Xia, Xinzhi Mu, Bo Pang, and Cewu Lu. Asynchronous interaction aggregation for action detection. In *ECCV*, 2020. 2, 8
- [41] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015. 1, 2
- [42] Du Tran, Heng Wang, Lorenzo Torresani, and Matt Feiszli. Video classification with channel-separated convolutional networks. *ICCV*, 2019. 1, 2, 7
- [43] 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. 1, 2, 3, 5, 6
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 4
- [45] Heng Wang, Du Tran, Lorenzo Torresani, and Matt Feiszli. Video modeling with correlation networks. In *CVPR*, 2020. 1, 2, 5, 6, 7, 8
- [46] 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. 2
- [47] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018. 1, 2, 4, 5, 6, 7, 8
- [48] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *ECCV*, 2018. 8
- [49] Jason Weston, Sumit Chopra, and Antoine Bordes. Memory networks. *arXiv preprint arXiv:1410.3916*, 2014. 2
- [50] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *CVPR*, 2019. 1, 2, 6, 8
- [51] Z. Wu, H. Li, C. Xiong, Y. G. Jiang, and L. S. Davis. A dynamic frame selection framework for fast video recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020. 2
- [52] Fanyi Xiao, Yong Jae Lee, Kristen Grauman, Jitendra Malik, and Christoph Feichtenhofer. Audiovisual slowfast networks for video recognition. *arXiv preprint arXiv:2001.08740*, 2020. 8
- [53] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning for video understanding. In *ECCV*, 2018. 1, 2, 4, 7, 8
- [54] 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. 2
- [55] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, 2018. 8
- [56] Linchao Zhu, Laura Sevilla-Lara, Du Tran, Matt Feiszli, Yi Yang, and Heng Wang. FASTER recurrent networks for efficient video classification. In *AAAI*, 2020. 1, 2
- [57] Mohammadreza Zolfaghari, Kamaljeet Singh, and Thomas Brox. ECO: Efficient convolutional network for online video understanding. In *ECCV*, 2018. 2, 8