



HAL
open science

ACDnet: An Action Detection network for real-time edge computing based on flow-guided feature approximation and memory aggregation

Yu Liu, Fan Yang, Dominique Ginhac

► To cite this version:

Yu Liu, Fan Yang, Dominique Ginhac. ACDnet: An Action Detection network for real-time edge computing based on flow-guided feature approximation and memory aggregation. *Pattern Recognition Letters*, 2021, 145, pp.118-126. 10.1016/j.patrec.2021.02.001 . hal-03153510

HAL Id: hal-03153510

<https://hal.science/hal-03153510>

Submitted on 26 Feb 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ACDnet: An Action Detection network for real-time edge computing based on flow-guided feature approximation and memory aggregation

Yu Liu**, Fan Yang, Dominique Ginjac

ImViA EA7535, Univ. Bourgogne Franche-Comté, Dijon 21078, France

ABSTRACT

Interpreting human actions requires understanding the spatial and temporal context of the scenes. State-of-the-art action detectors based on Convolutional Neural Network (CNN) have demonstrated remarkable results by adopting two-stream or 3D CNN architectures. However, these methods typically operate in a non-real-time, offline fashion due to system complexity to reason spatio-temporal information. Consequently, their high computational cost is not compliant with emerging real-world scenarios such as service robots or public surveillance where detection needs to take place at resource-limited edge devices. In this paper, we propose ACDnet, a compact action detection network targeting real-time edge computing which addresses both efficiency and accuracy. It intelligently exploits the temporal coherence between successive video frames to approximate their CNN features rather than naively extracting them. It also integrates memory feature aggregation from past video frames to enhance current detection stability, implicitly modeling long temporal cues over time. Experiments conducted on the public benchmark datasets UCF-24 and JHMDB-21 demonstrate that ACDnet, when integrated with the SSD detector, can robustly achieve detection well above real-time (75 FPS). At the same time, it retains reasonable accuracy (70.92 and 49.53 frame mAP) compared to other top-performing methods using far heavier configurations. Codes will be available at <https://github.com/dginjac/ACDnet>.

1. Introduction

In past years, human action detection has been an active area of research driven by numerous applications: autonomous vehicles, video search engines, and human-computer interaction, etc. As it aims not only to recognize actions of interest in a video, but also to localize each of them, action detection poses more challenges when compared to video classification. The task becomes even more difficult in practical applications when detection is to be performed in an online setting and at real-time speed. For instance, time-critical scenarios such as autonomous driving demand instant detection in order for machines to react immediately. Other use cases which seek for mobile or large-scale deployment, such as service robots and distributed unmanned surveillance, require detection or scene meta-data extraction at low-end edge devices. In general, edge devices (e.g., embedded systems) have limited computational power and are only compliant with resource-efficient detection algorithms.

Following the success of Convolutional Neural Network (CNN) in diverse computer vision tasks, modern action detectors are mainly based on CNN. In particular, fast object detectors have been widely adopted to spatially localize action instances at each frame (Singh et al. (2017), Zhao and Snoek (2019)). Naturally, effective temporal modeling plays an imperative role for identifying an action. To reason both spatial and temporal context, Simonyan and Zisserman (2014) pioneered the two-stream CNN framework which aggregates spatial and temporal cues from separate networks and input modalities (RGB and optical flow). Such an approach has motivated many state-of-the-art methods in the field of action recognition and detection. Alternatively, 3D CNN (Carreira and Zisserman (2017)) which performs spatio-temporal feature learning on stacked frames has also been increasingly explored to tackle video analysis tasks.

Despite recent advances in action detection, existing methods are inherently sub-optimal in two aspects. First, consecutive video frames exhibit high appearance similarity. Extracting frame features without taking into account this inter-frame similarity introduces redundancy. Moreover, the increased system complexity associated with employing two-stream or 3D CNN

**Corresponding author:

e-mail: yu_liu@etu.u-bourgogne.fr (Yu Liu),
fanyang@u-bourgogne.fr (Fan Yang), dominique.ginjac@ubfc.fr
 (Dominique Ginjac)

models is not proportionally reflected in the detection accuracy. In contrast, the above inevitably raises computational requirements associated with motion extraction and 3D convolution operation, prohibiting practical deployment on edge devices.

This work focuses on action detection solutions more pertinent to the criteria of realistic applications. To address the aforementioned limitations, we first exploit the temporal coherence among nearby video frames to enhance detection efficiency. This is embodied by performing feature approximation at the majority of frames in a video, mitigating re-extraction of similar features from neighboring frames. Furthermore, we hypothesize that a less expensive framework can effectively extract meaningful temporal contexts. Here, we adopt a multi-frame feature aggregation module, which recursively accumulates 2D spatial features over time to encapsulate long temporal cues. Such feature aggregation implicitly models temporal variations of actions and facilitates understanding degenerated frames with limited visual cues.

To the best of our knowledge, this is the first attempt applying feature approximation and aggregation techniques to achieve efficient action detection which can benefit resource-limited devices. To summarize, our contribution is three-fold:

- We propose an integrated detection framework, ACDnet, to address both detection efficiency and accuracy. It combines feature approximation and memory aggregation modules, leading to improvement in both aspects.
- Our generalized framework allows for smooth integration with state-of-the-art detectors. When incorporated with SSD (single shot detector), ACDnet could reason spatio-temporal context well over real-time, more appealing to resource-constrained devices.
- We conduct detailed studies in terms of accuracy, efficiency, robustness and qualitative analysis on public action datasets UCF-24 and JHMDB-21.

2. Related work

Recent advancements in action detection are largely led by building upon successful cases in object detection and action recognition. Here, we briefly review these relevant topics.

Object detection based on CNN methods can be grouped into two families. The two-stage approach such as Faster R-CNN by Ren et al. (2015) and R-FCN by Dai et al. (2016) first extracts potential object regions from images, on which it performs object classification and bounding box regression on features corresponding to each proposed location. Such a sequential pipeline imposes a bottleneck to real-time inference. Alternatively, single-stage detectors such as YOLO proposed by Redmon and Farhadi (2017), or SSD in Liu et al. (2016), remove the intermediate region proposal, directly achieving bounding box regression and classification in a single forward-pass. Bypassing the intermediate bottleneck enables real-time detection at the cost of minor accuracy drop.

A number of research focuses on video object detection instead of the image domain. Popular approaches such as Han

et al. (2016) and Kang et al. (2017) exploit videos' temporal consistency by associating detection boxes and scores from multiple frames. Similarly but on the feature level, Hetang et al. (2017) and Zhu et al. (2017a) aggregate multiple frame features to enhance detection accuracy. On the other hand, Zhu et al. (2017b) leverage the temporal redundancy among video frames to improve detection efficiency. Their framework propagates features from a sparse set of key frames to successive ones by motion to avoid re-extracting similar object features. In a similar spirit, Liu and Zhu (2018) propagate frame-level information across frames using a recurrent-convolutional architecture.

Action recognition is typically treated as a classification task on trimmed videos (Yao et al. (2019)). In addition to spatial features, reasoning temporal information across multiple frames is also crucial. Among different temporal modeling techniques, the two-stream architecture in Simonyan and Zisserman (2014) demonstrates state-of-art performance. Its framework consists of two feed-forward pathways, with one CNN learning spatial features from RGB stream and the other one learning motion features from optical flow stream. The two streams are trained and run inference independently to aggregate complemented features (Feichtenhofer et al. (2016)). Even though such a framework can exploit existing 2D CNN backbones, fine-grained optical flow is expensive to extract. Thus, flow images are typically pre-computed, which do not conform to the online workflow demanded in real-world scenarios.

Recently, 3D CNNs have been increasingly explored (Carrreira and Zisserman (2017), Li et al. (2019)) along with the release of large-scale action dataset Kinetics. They utilize 3D kernels to jointly perform spatio-temporal feature learning from stacked RGB frames, achieving comparable and even superior modeling capability than two-stream CNN. However, these models inherently suffer from higher number of parameters and computational cost than their 2D counterparts, making their deployment on resource-constrained devices impractical.

Efficient spatio-temporal modeling. To alleviate the high computational cost associated with flow extraction, several studies seek alternative motion representations that are easier to compute. These include feature-level displacement (Jiang et al. (2019), Sun et al. (2018b)), or simply taking the RGB difference between adjacent frames (Wang et al. (2016)). On the other hand, to reduce the complexity of 3D CNN, decoupled architectures such as P3D (Qiu et al. (2017)) and R(2+1)D (Tran et al. (2018)) have been studied. Alternatively, TSM proposed by Lin et al. (2019) handles temporal convolution as channel-shifting operators to fuse spatial features from different time steps. Their approach has demonstrated effectiveness on edge devices such as Jetson Nano and Galaxy Note8.

Spatio-temporal action detection simultaneously addresses action localization and classification in time and space. Leading approaches often leverage CNN object detectors as the core building block. The extension mainly consists of adopting the two-stream framework, fusing complementary detection results from both spatial and temporal stream to acquire frame-level detection, as demonstrated in Singh et al. (2017). For temporal localization, detection at each frame is then linked over time to construct action tubes (Peng and Schmid (2016)). Beyond

detection at the frame-level, Kalogeiton et al. (2017) and Li et al. (2020) adopt a clip-based approach, which exploits stacking multiple frame features to capture temporal cues on top of the two-stream architecture. In this case, actions are regressed and inferred directly on action cuboids.

Inspired by the latest adoption of 3D CNN in action recognition, more recent studies incorporate 3D CNN as the backbone (Yang et al. (2019), Sun et al. (2018a), Girdhar et al. (2019), Gu et al. (2018), Wei et al. (2019)). In addition, various ways of fusing spatial and temporal context have also been investigated. Besides aggregating at the detection level (e.g., union of detection results), others perform feature-level fusion. These include the use of 1×1 convolution (Ali and Taylor (2018)), attention model (Köpüklü et al. (2019)) or conditional normalization (Zhao and Snoek (2019)). Such approaches allow a part of CNN layers to adaptively learn from the fused features.

3. ACDnet

Our objective is to perform detection in an online manner for every incoming frame of a video. The proposed ACDnet which consists of the feature approximation and aggregation module, is summarized in Figure 1.

3.1. Feature approximation by motion guidance

Video content varies slowly over consecutive frames. This phenomenon is more so reflected in the corresponding CNN feature maps which capture high-level semantics. Intuitively, the shared appearances among neighboring frames can help to propagate essential information for a given task. The practice of propagation in Zhu et al. (2017b) has established success to enhance object detection efficiency in videos, which motivates our feature approximation module.

Within the approximation scheme, the heavier feature extraction sub-network, N_{feat} , only operates on a sparse set of key frames during inference. The features of successive non-key frames are obtained by spatially transforming those from their preceding key frames via two-channel flow fields. The workflow can be summarized by the following equations. Let $M_{i \rightarrow k}$ be the two-channel flow field capturing relative motion from the current frame I_i to its previous key frame I_k (horizontal and vertical direction). Then, feature approximation (also referred as feature propagation) is realized according to inverse warping:

$$F_i = W(F_k, M_{i \rightarrow k}) \quad (1)$$

where F_k is the key frame feature, and F_i is the newly warped feature corresponding to I_i . W denotes the inverse warping operation to sample the correct key frame features and assign them to the warped ones. Inverse warping is necessary to ensure every location p at the warped feature can be projected back to a point $p + \Delta p$ at the key frame feature, where $\Delta p = M_{i \rightarrow k}(p)$. Concretely, the warping operation W is performed as:

$$f_i^c(p_i) = \sum_{p_k} G(p_k, p_i + \Delta p) f_k^c(p_k) \quad (2)$$

In Equation 2, f_i^c and f_k^c denote the c^{th} channel of feature F_i and F_k , respectively; G denotes the bilinear interpolation kernel. Every location p_i in the warped feature map undergoes this warping scheme to sample features from key frames, independently for each feature channel c . The warping operation is much lighter compared to layers of convolution for feature extraction. Consequently, by applying feature approximation on a dense set of non-key frames, computation is greatly reduced.

Previous methods on action-based tasks typically acquire motion features from accurate optical flow using non-learning-based algorithms. However, computing flows in such a way imposes a bottleneck to real-time and online detection due to high consumption of time or requiring to pre-compute flow results. In contrast, ACDnet integrates a fast flow estimation sub-network, N_{flow} , to predict flow fields. In our case, optical flow serves to spatially transform CNN features; it does not need to capture fine-grained motion details and has the same height and width as the corresponding feature. Using such a learning-based flow estimator also allows it to be jointly trained with all other sub-networks specific to the task of action detection.

In detail, the flow sub-network take (I_k, I_i) as input, and generates a pair of motion field and position-wise scale map. Given that H , W , and C denote height, width and channel of F_k , then the flow field $M_{i \rightarrow k}$ is of size $H \times W \times 2$, and the scale map is $H \times W \times C$, whose dimension matches that of F_k to be warped. After the inverse warping described by Equation 1, the warped feature F_i is refined by multiplying the scale map in an element-wise way. Any F_k and F_i would be fed to the shared detection sub-network, N_{det} , to obtain final detection. This workflow is illustrated in Figure 1 (a) and (b).

3.2. Memory feature aggregation

Propagating features across frames reduces the computation cost associated with feature extraction. However, since most features are now approximated and heavily dependent of the quality of the precedent key frame features, we adopt a memory aggregation module as inspired by Hetang et al. (2017) to enhance the feature representation at key frames. Given incoming video frames, the core of memory aggregation is to reinforce features of a target frame by recursively incorporating supportive and discriminating context from the past. This allows implicit spatio-temporal modeling without explicitly extracting motion features. In addition, in cases when the current frame is deteriorated, an action can still be inferred with the supportive visual cues from memory. Figure 1 (c) gives an example when such memory aggregation could be useful.

Memory aggregation shares the same warping operation used for feature approximation. ACDnet takes a sparse and recursive approach to aggregate memory features only at key frames, due to similar appearances shared among nearby frames. Given two succeeding key frames I_{k1} and I_{k2} , where I_{k2} is the more recent one in time, memory aggregation follows Equation 3:

$$F_{k2_aggregated} = w_{k1} \otimes F'_{k1} + w_{k2} \otimes F_{k2} \quad (3)$$

where $F'_{k1} = W(F_{k1}, M_{k2 \rightarrow k1})$, is the warped feature of I_{k1} to spatially align its position with that of I_{k2} . The position-wise weights w_{k1} and w_{k2} both have the same height and width as

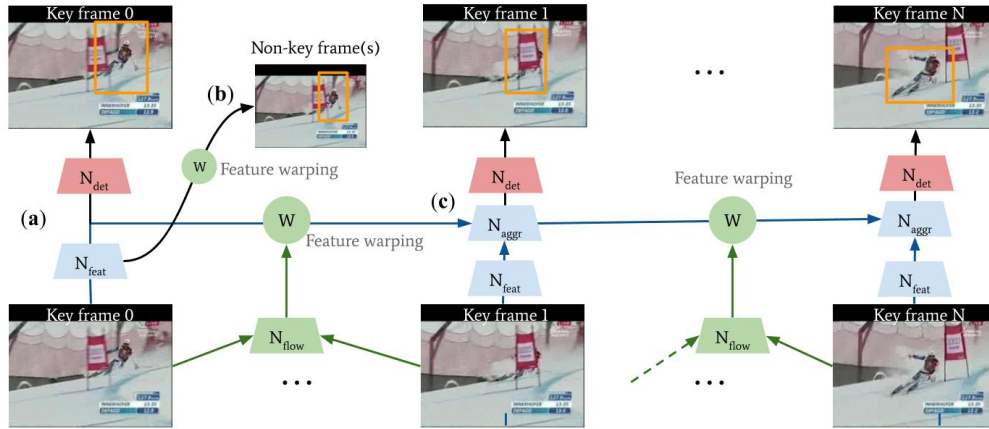


Fig. 1. Illustration of ACDnet inference pipeline. (a) At the initial frame, features are obtained from the feature extraction sub-network (N_{feat}). (b) For non-key frames (dense), the flow sub-network (N_{flow}) estimates a pair of flow field and position-wise scale map between the non-key frame and its preceding key frame. The resulted flow field is used to propagate appearance feature, which is then refined by the scale map via element-wise multiplication. (c) At key frames (sparse), new features are extracted. They are then aggregated with those from the past key frames (memory features) via N_{flow} and the aggregation sub-network (N_{agg}). The fused features will be used for detection (N_{det}) and also passed along as the updated memory.

F'_{k1} and F_{k2} . These weights are normalized and determine the importance of memory feature at each location p with respect to the target frame feature map ($w_{k1}(p) + w_{k2}(p) = 1$). All channels of a feature share the same spatial weights.

The weights w_{k1} and w_{k2} are adaptively calculated based on the similarity of memory and target features. We estimate feature similarity by first projecting them into an embedding space via a few convolution layers, and then computing the cosine similarity between the embedded features. Finally at the current key frame, the weighted sum of the memory and current features will be fed to the detection sub-network and passed along as the new memory.

3.3. Training procedure

ACDnet follows a three-frame training scheme, as depicted in Figure 2. From each training mini-batch, frame I_i and two precedent video frames (I_k and I_{mem}) are selected, whose features simulate key frame and memory features respectively. The offset between I_i and I_k is a random number from 0 to $T_{k \rightarrow i}$, and the offset between I_k and I_{mem} is fixed at $T_{mem \rightarrow k}$.

The feature maps F_{mem} and F_k are first extracted from I_{mem} and I_k respectively. Two sets of flow fields, namely, the relative motion between $I_k - I_{mem}$, and $I_i - I_k$ are also estimated. The former flow is used to propagate F_{mem} to F_k to simulate the occurrence of memory feature aggregation following Equation 3. Finally, the fused feature is warped with the second flow (simulating feature approximation) following Equation 1, which will be the final feature map for N_{det} . Under this training mode, only the groundtruth of I_i is needed to determine losses, which are back-propagated to update all sub-networks.

3.4. Adaptation for multi-feature scale detector

Workflows of feature approximation and aggregation are generic for video-based tasks. ACDnet further employs SSD, an one-stage detector to fulfill the objective of high-speed action detection potentially for embedded vision systems. In particular, the SSD300 model is chosen due to its superior speed.




Fig. 2. Training procedure. Each mini-batch consists of three frames (I_{mem} , I_k , and I_i) and the groundtruth of I_i .

In SSD, a set of auxiliary convolutional layers are progressively added after the base network (e.g., VGG16 in a standard SSD) to extract features at multiple scales. This creates multiple feature maps where the detector makes prediction for objects of various sizes. Consequently, adopting the described framework in SSD requires feature approximation and memory aggregation to be handled for features at all scales.

To enable multi-level feature approximation, we duplicate N_{flow} 's flow prediction layer into several branches. The number of branches matches that of the feature maps; the branches' outputs are also progressively resized via average pooling according to sizes of SSD's feature maps. Then, each branch reconstructs a pair of flow field and scale map in accordance with the dimension of SSD feature (refer to Figure 3). To cope with multi-level feature approximation and aggregation, Equation 1 and 3 are also generalized to take place at each feature level independently. Note that the standard SSD300 applies detection at six feature scales. Nevertheless, we only use the first five of them, as the dimension of the last feature map becomes a 1D vector resulting from progressive resizing, which is not feasible for feature approximation governed by 2D spatial warping.

4. Experimental results

4.1. Experimental Setup

Dataset. Our proposed methods are evaluated on two popular action datasets: **UCF-24** and **JHMDB-21**. The former one




Fig. 3. Flow estimation sub-network adapted for multi-scale feature approximation and aggregation. The depicted design corresponds to the architecture of SSD300 and FlowNet.

released by Soomro et al. (2012) is composed of 3207 sports videos of 24 action classes. Following previous work, we use 2290 of these video clips for training. The latter collected by Jhuang et al. (2013) consists of 928 short videos divided into three splits, with 21 action categories in daily life. Each video is trimmed and has a single action instance. We report our experimental results on the average of three splits for this dataset.

Network architectures. ACDnet incorporates the following sub-networks: SSD300 (with VGG16 backbone), FlowNet (Dosovitskiy et al. (2015)) and feature embedding. Feature embedding contains five branches for measuring feature similarity at five different scales. Each embedding branch has a bottleneck design of three 1×1 convolution layers interleaved with ReLU non-linearity, where the number of filters corresponds to the number of channels at each feature level l : $feat_{channel}^l/2$, $feat_{channel}^l/2$ and $feat_{channel}^l \times 2$, respectively.

FlowNet is modified to also generate five sets of flow fields and position-wise scale maps, each pair being used for warping and refining designated features. We initialize the weights of the first two branches of flow generation layers using FlowNet’s pre-trained weights. Considering the later three flow outputs are spatially much smaller than that of the original FlowNet, we randomly initialize the weights of those branches.

Training. Images are resized to 300×300 for training and inference. Training is conducted by stochastic gradient descent. To address data imbalance among different actions, from each training video clip of UCF-24, 15 frames spanning the whole video are evenly sampled as the training set. Since video clips of JHMDB-21 are generally short (≤ 40), we evenly sample 10 frames from each clip for training. Specifically, both $T_{mem \rightarrow k}$ and $T_{k \rightarrow i}$ are set to 10 during training. These chosen values correspond to the key frame interval used during inference, which is also fixed at 10 in our experiments unless specified.

We apply different hyperparameters on the two datasets. UCF-24 is trained for 100K iterations; the learning rate is initialized as 0.0005 and reduced by a factor of 0.1 after the 80Kth and 90Kth iterations. Weights of VGG16’s first two convolution blocks are frozen. For JHMDB-21, due to its smaller training and testing size, we observe that detection accuracy tends to fluctuate significantly between successive epochs. Hence, we empirically train this dataset for 20K iterations with learning rate initialized as 0.0004 and reduced by a factor of 0.5 after the 8Kth and 16Kth iteration. During its training, the first three

Table 1. F-mAP results for different configurations.

Table 1. F-mAP results for different configurations.					
ACDnet					
SSD	✓	✓	✓	✓	✓
FA		✓	✓	✓	✓
Scale map			✓		✓
MA				✓	✓
F-mAP					
UCF-24	67.32	65.84	67.23	68.06	70.92
JHMDB-21	47.90	46.65	46.69	49.37	49.53

convolution blocks of VGG16 are frozen. In addition, all layers of FlowNet until the flow generation layers (the five branches at the end of our modified model) are also frozen to further reduce the risk of overfitting.

All sub-networks are trained jointly (also evaluated) on an NVIDIA Quadro P6000 GPU using a training batch size of 8. For the rest of hyperparameters and data augmentation methods, we follow the same setup as the original SSD by Liu et al. (2016). The weights of VGG16 and FlowNet are pre-trained using ImageNet and the Flying Chair dataset respectively.

4.2. Ablation study

Our proposed architecture has been evaluated in terms of accuracy, efficiency and robustness over several network configurations. The standard frame-level mean average precision (F-mAP) and frame-per-second (FPS) have been used as the evaluation metrics. Specifically, FPS is measured based on the complete detection pipeline, including data loading and model inference using batch size of 1. The Intersection-over-Union threshold is set to 0.5 throughout all experiments. For brevity, we refer to feature approximation and memory feature aggregation as **FA** and **MA** respectively when presenting their results.

Accuracy. F-mAP results of different configurations are reported in Table 1. From both datasets, we observe a decrease of accuracy when only feature approximation is included. However, the accuracy drop can be compensated by the addition of memory aggregation, which exceeds the accuracy of the stand-alone SSD. Figure 4 shows some examples of how the memory aggregation module benefits detection. Overall, we remark that aggregating multipleframe features over time, even in a sparse manner, improves models’ abilities to more confidently discriminate among different actions.

We also examine the effect of separate branches of position-wise scale maps designed for refining visual features. Our results indicate that such refinement mildly improves detection accuracy. The scale maps serve as implicit attention maps which reinforce feature responses associated with moving actors (elaborated in Figure 5).

Even though similar result patterns can be seen from both datasets, the benefit of memory aggregation appears less prominent in JHMDB-21. This could result from the fact that each video clip in JHMDB-21 is much shorter (40 frames or less). As MA is performed sparsely at every 10th frame, its impact is limited to 2-3 aggregation per clip. Furthermore, we observe that motions in several JHMDB-21 clips are relatively small.




Fig. 4. Examples where ACDnet (FA, MA) improves the baseline SSD. Green / Red boxes correspond to correct / incorrect detection respectively.




Fig. 5. Position-wise scale maps produced by our modified FlowNet. The scale map (bottom row) only reinforces activation (top row) associated with the actor by up-scaling, without altering activation in other feature regions.

In these clips, key frames far apart still appear fairly identical, limiting additional visual cues to be propagated.

Efficiency is evaluated on UCF-24 by simultaneously inspecting accuracy, run time and number of parameters of various configurations. Here, we assume the use of scale map refinement if applicable. As shown in Table 2, ACDnet (SSD, FA, MA) outperforms the stand-alone SSD in both speed and accuracy. This suggests it is relevant to handle inter-frame redundancy, and that long-range memory fusion is effective for collecting more discriminating features. Regarding the number of required parameters, the increase in ACDnet (SSD, FA) from stand-alone SSD is associated with the addition of FlowNet, which can be replaced by much lighter architectures in the future. Likewise, the increase with the addition of MA module corresponds to the extra embedding layers for measuring feature similarity at various scales. In terms of run time, the speed drop with MA is incurred by the additional operations at key frames (except for the first one), where flow estimation, feature extraction, similarity measure and aggregation all take place.

To examine how our generic architecture performs on a different detection framework, we conduct the same experiments while incorporating ACDnet with R-FCN, a state-of-the-art two-stage detector. The run time improvement brought by feature approximation is more significant with R-FCN, due to it using a much deeper backbone for feature extraction. The number of additional parameter needed to carry out memory aggregation is less too for R-FCN, as it is designed to perform prediction on a single-scale feature (needing only one branch for the

Table 2. Performance of different configurations on UCF-24.

	F-mAP	FPS	# params
SSD	67.32	70	26.8M
ACDnet (SSD, FA)	67.23	85	50.8M
ACDnet (SSD, FA, MA)	70.92	75	57M
R-FCN	68.2	15	60M
ACDnet (R-FCN, FA)	66.19	33.5	85.7M
ACDnet (R-FCN, FA, MA)	68.31	32	89.6M




Fig. 6. F-mAP under varied key frame intervals.

embedding and flow sub-networks). Overall, when taking into account run time, memory consumption and obtained accuracy together, our results still strongly favor the SSD-based ACDnet.

Robustness. Concerning the robustness of our models trained with a fixed duration $T_{mem \rightarrow k}$ and $T_{k \rightarrow i}$ (at 10), we evaluate their performances under various key frame intervals (k) during inference. Figure 6 displays F-mAP results on both datasets, with k ranging from 2 to 20. Both models (with and without MA) express an overall steady drop in accuracy on the two datasets as k increases. This is reasonable, as the ability of flow fields to correctly encode pixel correspondence diminishes under large motions. However, even when k is large, ACDnet with MA still retains decent accuracy which outperforms or is comparable with the best cases of the other configurations.

Run time is also inspected under the same setting, as shown in Figure 7. It can be observed that ACDnet (FA, MA) exceeds the speed of SSD starting around $k = 8$, while the FA-only model is consistently faster. Larger key frame intervals intuitively should lead to further speed boost, as higher ratio of features are approximated. Interestingly, we observe that this pattern is neatly presented when $k \leq 10$. After that, the run time of the examined models begin to saturate. This phenomenon is associated with two factors. On the one hand, as key frame interval increases, the ratio between approximated and real features become less significant. On the other hand, larger key frame intervals introduce more motion which could compromise the quality of approximated features. This results in an increase of low-confident predictions, which take longer for SSD's non-maximum suppression to filter.




Fig. 7. FPS under varied key frame intervals.

4.3. Comparison with state-of-the-art

We compare the complete ACDnet (with FA and MA) against state-of-the-art methods in Table 3. Since our proposed framework targets lightweight action inference for realistic deployment rather than solely obtaining superior accuracy, only top-performing works which take into account both accuracy and run time are considered for fair comparison. With this in mind, recent research such as the works of Wei et al. (2019) and Gu et al. (2018) demonstrate impressive accuracy but are excluded from our comparison, as they utilize heavier configurations and omit speed analysis. Alongside performances, comprehensive summary of each method’s backbone is also reported for clearer comparison. It should be noted that methods such as ACT, STEP and MOC perform clip-based detection. In other words, they take clips of multiple RGB frames with the support of stacked flow images at once (e.g., five flow for each RGB), predicting action tubelets spanning these RGB frames. In contrast, methods such as YOWO gather supportive contextual cues from multiple frames to augment the target one. These particular attributes are summarized Table 3 column 4.

As shown in Table 3, ACDnet outperforms the others in terms of run time. This is ascribed to the feature approximation module and our less complex architecture overall. The other methods either adopt two-stream or 3D CNN architectures to capture complemented spatial and temporal features, which raise computation. In addition, preparation of accurate flow using Brox (Brox et al. (2004)) or FlowNet2 (Ilg et al. (2017)) is particularly expensive; as a result, all methods employing a second flow stream do not take into account optical flow acquisition when measuring run time (except ROAD using a fast flow estimator by Kroeger et al. (2016)). In contrast, flow generation in ACDnet is fast and can be carried out in an online setting as it does not aim to encode fine-grained motion features.

In terms of accuracy, ACDnet retains competitive performance on UCF-24. On the other hand, its performance on JHMDB-21 is less impressive compared to the other methods. As opposed to UCF-24, whose classes of sports activities are visually more distinctive, we observe that JHMDB-21 contains more classes that share similar visual context (for example, *Sit* v.s. *Stand*, and *Run* v.s. *Walk*, etc.). Figure 8 demonstrates



Fig. 8. Examples of false detection in JHMDB-21. (a) Correct action: *Jump*. (b) Correct action: *Sit*. (c). Correct action: *Stand*; ACDnet incorrectly predicts two actions (*Stand* and *Run*).

a few falsely detected examples by our model which result in lower F-mAP in JHMDB-21. Such ambiguous visual context is challenging even for human to confidently infer the correct action unless viewing consecutive frames at once. As shown in column 4 of Table 3, ACDnet applies detection on frames far fewer than other methods, which limits its ability to model detailed variations of visual cues over time. In addition, JHMDB-21 consists of short clips for which sparse memory aggregation can only take place minimally. The above factors result in ACDnet’s less satisfactory accuracy on JHMDB-21. This visual ambiguity could generally be mitigated when examining more frames at once, as demonstrated by all clip-based methods.

Similarly, 3D-CNN-based method such as YOWO also proves effective to learn spatio-temporal features when taking 16 consecutive frames. However, such an approach inevitably raises the computation time; not only from model inference, but also data loading, which is excluded in their reported speed performance. Furthermore, ACDnet achieves comparable accuracy when YOWO employs lighter 3D CNN variants, implying the necessity of deeper models to effectively reason temporal context. In conclusion, our experimental results verify ACDnet’s competitive capability to efficiently infer actions with strong visual cues, but its sparse spatio-temporal modeling scheme does not capture temporal cues as effectively as the more expensive two-stream / 3D CNN. On the other hand, ACDnet is compact and can achieve inference speed far beyond real-time requirement. This not only permits more seamless deployment potentially on resource-constrained devices, but also can afford to further adopt a clip-based framework or lightweight 3D CNN to improve its accuracy.

5. Conclusions and Future works

In this paper, we present ACDnet, a compact action detection network with real-time capability. By exploiting temporal coherence among video frames, it utilizes feature approximation on frames with similar visual appearances, which significantly improves detection efficiency. Additionally, a memory aggregation module is introduced to fuse multi-frame features, enhancing detection stability and accuracy. The combination of the two modules and SSD detector implicitly reasons temporal context in an inexpensive manner. ACDnet demonstrates real-time detection (up to 75 FPS) on public benchmarks while retaining decent accuracy against other best performers at far less complex settings, making it more appealing to edge device deployment in practical applications. Our future works include further investigation in cost-effective architectures for

Table 3. State-of-the-art comparison. *For any key frame, ACDnet (FA, MA) fuses the accumulated key frame feature from the past with the current one (considered 2RGB implicitly). For any non-key frame, its feature is approximated based on the preceding key frame feature (considered 1RGB).

Method	+2-stream Flow	+3D CNN	#frames:#det.	UCF24	JHMDB21	FPS
ROAD, Singh et al. (2017)	Kroeger	✗	(1RGB+1FL):1	65.66	–	28
ACT, Kalogeiton et al. (2017)	Brox	✗	(6RGB+30FL):6	69.5	65.7	25
TS-YOLO, Ali and Taylor (2018)	FlowNet2-SD	✗	(1RGB+1FL):1	71.67	–	25
STEP, Yang et al. (2019)	Brox	3× 3D conv	(6RGB+30FL):6	75	–	21
YOWO(a), Köpüklü et al. (2019)	✗	3D-ResNet101	16RGB:1	80.4	74.4	34
YOWO(b)	✗	3DShuffleNetV2	16RGB:1	71.4	55.3	–
YOWO(c)	✗	3DMobileNetV2	16RGB:1	66.6	52.5	–
MOC, Li et al. (2020)	Brox	✗	(7RGB+35FL):7	78	70.8	25
ACDnet (FA, MA)	✗	✗	2(1)*RGB:1	70.92	49.53	75

spatio-temporal modeling and performing temporal localization. For a fully integrated and resource-efficient vision system, lightweight alternatives of the current sub-networks will be explored, and we will precisely customize candidate solutions for embedding them on edge devices such as NVIDIA Xavier GPU.

Acknowledgments

This work was supported by the H2020 ITN project ACHIEVE (H2020-MSCA-ITN-2017: agreement no. 765866).

References

- Ali, A., Taylor, G.W., 2018. Real-time end-to-end action detection with two-stream networks, in: IEEE CRV, pp. 31–38.
- Brox, T., Bruhn, A., Papenberger, N., Weickert, J., 2004. High accuracy optical flow estimation based on a theory for warping, in: ECCV, Springer. pp. 25–36.
- Carreira, J., Zisserman, A., 2017. Quo vadis, action recognition? a new model and the kinetics dataset, in: IEEE CVPR, pp. 6299–6308.
- Dai, J., Li, Y., He, K., Sun, J., 2016. R-fcn: Object detection via region-based fully convolutional networks, in: NIPS, pp. 379–387.
- Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T., 2015. FlowNet: Learning optical flow with convolutional networks, in: IEEE ICCV, pp. 2758–2766.
- Feichtenhofer, C., Pinz, A., Zisserman, A., 2016. Convolutional two-stream network fusion for video action recognition, in: IEEE CVPR, pp. 1933–1941.
- Girdhar, R., Carreira, J., Doersch, C., Zisserman, A., 2019. Video action transformer network, in: IEEE CVPR, pp. 244–253.
- Gu, C., Sun, C., Ross, D.A., Vondrick, C., Pantofaru, C., Li, Y., Vijayanarasimhan, S., Toderici, G., Ricco, S., Sukthankar, R., et al., 2018. Ava: A video dataset of spatio-temporally localized atomic visual actions, in: IEEE CVPR, pp. 6047–6056.
- Han, W., Khorrani, P., Paine, T.L., Ramachandran, P., Babaeizadeh, M., Shi, H., Li, J., Yan, S., Huang, T.S., 2016. Seq-nms for video object detection. arXiv preprint arXiv:1602.08465 .
- Hetang, C., Qin, H., Liu, S., Yan, J., 2017. Impression network for video object detection. arXiv preprint arXiv:1712.05896 .
- Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T., 2017. FlowNet 2.0: Evolution of optical flow estimation with deep networks, in: IEEE CVPR, pp. 2462–2470.
- Jhuang, H., Gall, J., Zuffi, S., Schmid, C., Black, M.J., 2013. Towards understanding action recognition, in: IEEE ICCV, pp. 3192–3199.
- Jiang, B., Wang, M., Gan, W., Wu, W., Yan, J., 2019. Stm: Spatiotemporal and motion encoding for action recognition, in: IEEE ICCV, pp. 2000–2009.
- Kalogeiton, V., Weinzaepfel, P., Ferrari, V., Schmid, C., 2017. Action tubelet detector for spatio-temporal action localization, in: IEEE ICCV, pp. 4405–4413.
- Kang, K., Li, H., Yan, J., Zeng, X., Yang, B., Xiao, T., Zhang, C., Wang, Z., Wang, R., Wang, X., et al., 2017. T-cnn: Tubelets with convolutional neural networks for object detection from videos. IEEE TCSVT 28, 2896–2907.
- Köpüklü, O., Wei, X., Rigoll, G., 2019. You only watch once: A unified cnn architecture for real-time spatiotemporal action localization. arXiv preprint arXiv:1911.06644 .
- Kroeger, T., Timofte, R., Dai, D., Van Gool, L., 2016. Fast optical flow using dense inverse search, in: ECCV, Springer. pp. 471–488.
- Li, Y., Miao, Q., Tian, K., Fan, Y., Xu, X., Ma, Z., Song, J., 2019. Large-scale gesture recognition with a fusion of rgb-d data based on optical flow and the c3d model. PRL 119, 187–194.
- Li, Y., Wang, Z., Wang, L., Wu, G., 2020. Actions as moving points. arXiv preprint arXiv:2001.04608 .
- Lin, J., Gan, C., Han, S., 2019. Tsm: Temporal shift module for efficient video understanding, in: IEEE ICCV, pp. 7083–7093.
- Liu, M., Zhu, M., 2018. Mobile video object detection with temporally-aware feature maps, in: IEEE CVPR, pp. 5686–5695.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C., 2016. Ssd: Single shot multibox detector, in: ECCV, Springer. pp. 21–37.
- Peng, X., Schmid, C., 2016. Multi-region two-stream r-cnn for action detection, in: ECCV, Springer. pp. 744–759.
- Qiu, Z., Yao, T., Mei, T., 2017. Learning spatio-temporal representation with pseudo-3d residual networks, in: IEEE ICCV, pp. 5533–5541.
- Redmon, J., Farhadi, A., 2017. Yolo9000: better, faster, stronger, in: IEEE CVPR, pp. 7263–7271.
- Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster r-cnn: Towards real-time object detection with region proposal networks, in: NIPS, pp. 91–99.
- Simonyan, K., Zisserman, A., 2014. Two-stream convolutional networks for action recognition in videos, in: NIPS, pp. 568–576.
- Singh, G., Saha, S., Sapienza, M., Torr, P.H., Cuzzolin, F., 2017. Online real-time multiple spatiotemporal action localisation and prediction, in: IEEE ICCV, pp. 3637–3646.
- Soomro, K., Zamir, A.R., Shah, M., 2012. Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402 .
- Sun, C., Shrivastava, A., Vondrick, C., Murphy, K., Sukthankar, R., Schmid, C., 2018a. Actor-centric relation network, in: ECCV, pp. 318–334.
- Sun, S., Kuang, Z., Sheng, L., Ouyang, W., Zhang, W., 2018b. Optical flow guided feature: A fast and robust motion representation for video action recognition, in: IEEE CVPR, pp. 1390–1399.
- Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M., 2018. A closer look at spatiotemporal convolutions for action recognition, in: IEEE CVPR, pp. 6450–6459.
- Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Van Gool, L., 2016. Temporal segment networks: Towards good practices for deep action recognition, in: ECCV, Springer. pp. 20–36.
- Wei, J., Wang, H., Yi, Y., Li, Q., Huang, D., 2019. P3d-ctn: Pseudo-3d convolutional tube network for spatio-temporal action detection in videos, in: IEEE ICIP, pp. 300–304.
- Yang, X., Yang, X., Liu, M.Y., Xiao, F., Davis, L.S., Kautz, J., 2019. Step: Spatio-temporal progressive learning for video action detection, in: IEEE CVPR, pp. 264–272.
- Yao, G., Lei, T., Zhong, J., 2019. A review of convolutional-neural-network-based action recognition. PRL 118, 14–22.
- Zhao, J., Snoek, C.G., 2019. Dance with flow: Two-in-one stream action detection, in: IEEE CVPR, pp. 9935–9944.
- Zhu, X., Wang, Y., Dai, J., Yuan, L., Wei, Y., 2017a. Flow-guided feature aggregation for video object detection, in: IEEE ICCV, pp. 408–417.
- Zhu, X., Xiong, Y., Dai, J., Yuan, L., Wei, Y., 2017b. Deep feature flow for

video recognition, in: IEEE CVPR, pp. 2349–2358.