

# Fast Object Detection in Compressed Video

Shiyao Wang<sup>1,2</sup> \* Hongchao Lu<sup>1</sup> Zhidong Deng<sup>1</sup>

Department of Computer Science and Technology, Tsinghua University<sup>1</sup>  
Alibaba Group<sup>2</sup>

wangshy31@gmail.com luhc15@mails.tsinghua.edu.cn michael@tsinghua.edu.cn

## Abstract

Object detection in videos has drawn increasing attention since it is more practical in real scenarios. Most of the deep learning methods use CNNs to process each decoded frame in a video stream individually. However, the free of charge yet valuable motion information already embedded in the video compression format is usually overlooked. In this paper, we propose a fast object detection method by taking advantage of this with a novel Motion aided Memory Network (MMNet). The MMNet has two major advantages: 1) It significantly accelerates the procedure of feature extraction for compressed videos. It only need to run a complete recognition network for I-frames, i.e. a few reference frames in a video, and it produces the features for the following P frames (predictive frames) with a light weight memory network, which runs fast; 2) Unlike existing methods that establish an additional network to model motion of frames, we take full advantage of both motion vectors and residual errors that are freely available in video streams. To our best knowledge, the MMNet is the first work that investigates a deep convolutional detector on compressed videos. Our method is evaluated on the large-scale ImageNet VID dataset, and the results show that it is  $3\times$  times faster than single image detector R-FCN and  $10\times$  times faster than high-performance detector MANet at a minor accuracy loss.

## 1. Introduction

Video is viewed as one of the next frontiers in computer vision since many real-world data sources are video based, ranging from visual surveillance [3], human-computer interaction [30] to autonomous driving [46]. In the past five years, deep learning methods have made historic progress in still image analysis [37, 41, 14, 39, 55]. Novel CNN

\*The work was done when Shiyao Wang was at Tsinghua University.

<sup>1</sup>State Key Laboratory of Intelligent Technology and Systems, THUAI, BNRist, Center for Intelligent Connected Vehicles and Transportation, Tsinghua University, Beijing, China.

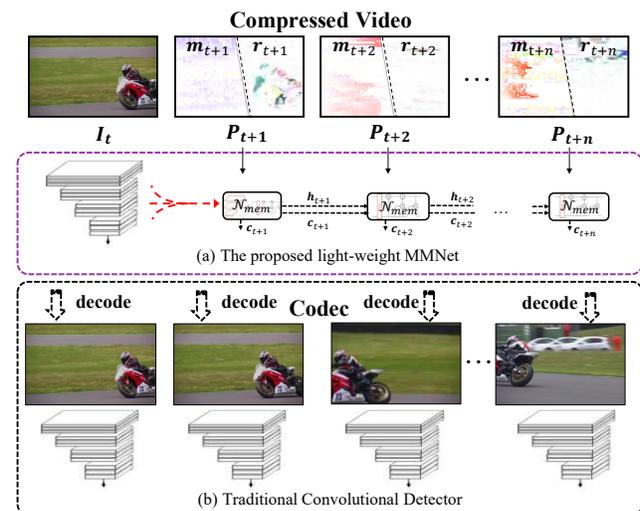


Figure 1. (a) The proposed light-weight MMNet accelerates the CNN inference by using compressed video directly, while (b) most previous methods for video analysis use heavy computational networks to extract features frame by frame.

based frameworks have been proposed for single image object detection, including Faster R-CNN [33], R-FCN [7], SSD [29], YOLO [32] and FPN [26]. Although there has great success in static image object detection, it still remains a challenging problem for detection in videos. Since frames may suffer from imaging-related degradations, most previous works [53, 44, 48, 21, 22] focus on improving frame-wise detection results. They extract features of the dense frames by applying existing image recognition networks (e.g., ResNet[14]) individually (see the bottom line in Figure 1), and leverage temporal coherence by feature aggregation or bounding box rescoring. Although these methods improve final performance, using CNNs to process the dense frames of videos is computationally expensive while it becomes unaffordable as the video goes longer.

In order to reduce the redundant computation, [54, 52] propose methods that run an expensive feature extractor

only on sparse keyframes and then propagate the resulting deep features to other frames. The key idea for feature propagation among frames is to calculate pixel-wise displacements through FlowNet[8]. However, it pays extra time for calculating displacements since the FlowNet still composes dozens of convolutional layers. They treat a video as a sequence of independent images and ignore the fact that a video is generally stored and transmitted in a compressed data format. The codecs split a video into I-frames (intra-coded frames) and P/B-frames (predictive frames). An I-frame is a complete image while a P/B-frame only holds the changes compared to the reference frame. For example, the encoder stores the object’s movements  $m_{t+k}$  and residual errors  $r_{t+k}$  when it moves across a stationary background (see the top line in Figure 1). So, consecutive frames are highly correlated, and the changes are already encoded in a video stream. Treating them as a sequence of still images and exploiting different techniques to retrieve motion cues seem time-consuming and cumbersome.

In this paper, we propose a fast and accurate object detection method for compressed videos called Motion-aided Memory Network with pyramidal feature attention (MM-Net). For a group of successive pictures (GOP) in a video, it runs the complete recognition network for I-frames, while a light-weight memory is developed to produce features for the following P-frames. The proposed MMNet receives the features of the preceding I-frame as input and fast predicts the following features by using motion vectors and residual errors in video streams. Moreover, different from the previous work that only propagates high-level features, the proposed memory network composes pyramidal features which enable the model to detect objects across multiple scales. In summary, the contributions of this paper include:

- We explore inherent motion signals and residual errors in codecs to align and refine features. Note that the signals retain necessary motion cues and are freely available.
- We propose a pyramidal feature attention that enables the memory network to propagate features from multiple scales. It helps to detect objects across different scales.
- We evaluate the proposed model on the large-scale ImageNet VID dataset [35] and present memory visualization for further analysis. Our model achieves significant speedup at a minor accuracy loss.

## 2. Related work

### 2.1. Object detection

#### 2.1.1 Object detection from still images

State-of-the-art methods for general object detection consist of feature networks [25, 37, 40, 14, 39, 18, 6] and detection networks [11, 10, 13, 33, 7, 26, 36, 27]. [11] is a typical proposal based detector which uses extracted proposals [43]. Faster R-CNN [33] further integrates proposal generation

step into CNNs. R-FCN [7] has comparable performance and higher speed compared to Faster R-CNN. We use R-FCN as our baseline and its computation speed is further improved for video object detection.

#### 2.1.2 Object detection in videos

One of the main-stream methods is based on per-frame complete detection and improves the detection quality by leveraging temporal coherence. And the other tries to speed up the computation by using temporal redundancy.

**For high performance**, [53, 21, 44, 48, 51, 2] propose end-to-end learning models to enhance per-frame features. [53, 44, 51] adopt FlowNet [8] to align and aggregate features. [21] provides a novel tubelet proposal network to efficiently generate spatiotemporal proposals. [48] computes the correlation between neighboring frames and introduces a memory module to aggregate their features. [2] uses deformable convolutions across time to align the features from the adjacent frames. [2, 22, 9] are based on detected bounding boxes rather than feature-level aggregation. [12, 23, 22] propose mapping strategies of linking still image detections to cross-frame box sequences in order to boost scores of weaker detections. D&T [9] is the first work to joint learn ROI tracker along with detector and the tracker is also exploited to link the cross-frame boxes. All of these mentioned works achieve high detection performance but they use a computationally expensive network to generate per-frame features.

**For fast inference**, [54] utilizes optical flow network for calculating pixel-level correspondence and propagating deep feature maps from keyframes to other frames. The flow estimation and feature propagation are faster than feature networks. Thus, the significant speedup is achieved.[5] introduces temporal propagation on box-level. They first produce bounding boxes on keyframes, and then generate boxes of other frames through a coarse-to-fine network. [28] propagates feature maps across frames via a convolutional LSTM. They only use appearance features without explicitly capturing motion cues. Although their model is faster than existing methods, the performance is much degraded. [45, 16, 19] also focus on model accelerating. They aim to build light-weight deep neural networks which are unrelated to specific tasks.

### 2.2. Deep learning model on compressed videos

H.264/MPEG-4 Part 10, Advanced Video Coding [38] is one of the most commonly used formats for recording, compression and distribution of videos. It is a block-oriented motion-compensation-based video compression standard [34]. To our knowledge, only a few prior works applied deep models directly on compressed videos. [24, 42] utilize signals from compressed video to produce non-deep fea-

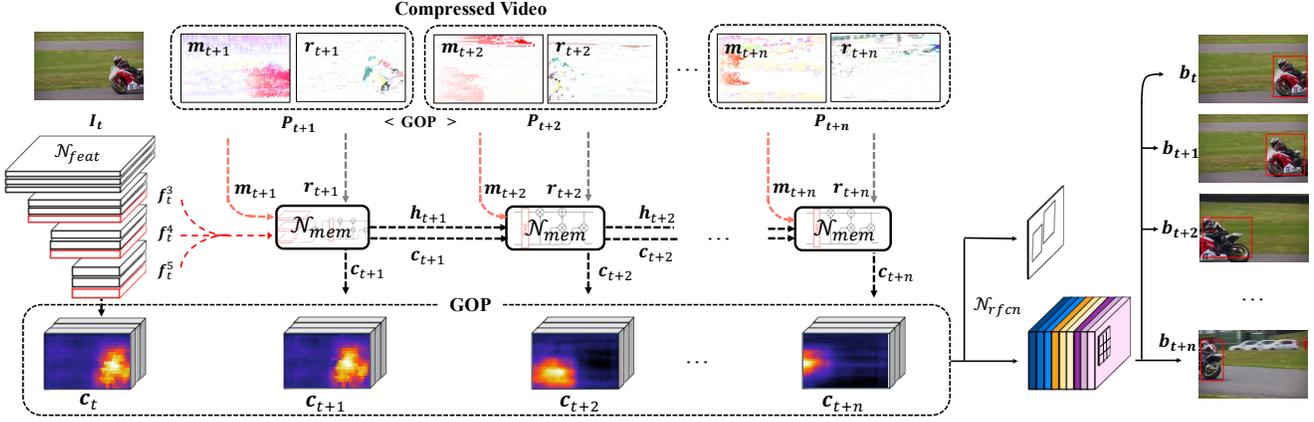


Figure 2. The overall framework of the proposed MMNet with pyramidal feature attention. The feature extractor  $\mathcal{N}_{feat}$  only runs on reference frame  $I_t$ , and the other features of  $c_{t+k}$  are generated by the memory network  $\mathcal{N}_{mem}$ . Motion vectors  $m_{t+k}$  and residual errors  $r_{t+k}$  are fed into the memory network so as to provide motion cues. Finally, all the features in one GOP (group of pictures) are aggregated to the detection network  $\mathcal{N}_{rfcn}$ , producing bounding boxes simultaneously.

tures. [47] resembles our model the most and they aim to improve both speed and performance on video action recognition which focuses on producing video-level features. But the video object detection needs to produce per-frame bounding boxes that has per-frame feature quality requirements.

### 3. Method

#### 3.1. Overview

The proposed motion-aided memory network with pyramidal feature attention is presented in Figure 2.

For the input video, we use H.264 baseline profile as illustration since these compression techniques that leverage consecutive frames are usually similar. H.264 baseline profile contains two types of frames: I- and P-frames. An I-frame (denoted as  $I_t \in \mathbb{R}^{h \times w \times 3}$ ) is a complete image.  $h$  and  $w$  are the height and width. P-frames are also known as delta-frames, denoted as  $P_{t+k}$ . They can be reconstructed by using the stored offsets, called motion vectors  $m_{t+k}$  and residual errors  $r_{t+k}$ . Detailed illustration of extracting  $m_{t+k}$  and  $r_{t+k}$  is presented in Section 3.3.1. In Figure 2, we show a typical GOP on the top line, denoted as  $\{I_t, P_{t+1}, \dots, P_{t+k}, \dots, P_{t+n}\}$ .

For the core modules, there are three networks: feature extractor, memory network and detection network, indicated as  $\mathcal{N}_{feat}$ ,  $\mathcal{N}_{mem}$  and  $\mathcal{N}_{rfcn}$ , respectively. The I-frame  $I_t$  is fed to the  $\mathcal{N}_{feat}$  in order to generate pyramidal features  $f_t^l \in \mathbb{R}^{h^l \times w^l \times c^l}$ .  $l$  is the index of multiple stages in a network and  $w^l$ ,  $h^l$  and  $c^l$  are the corresponding width, height and channel numbers. They are sent to the memory network  $\mathcal{N}_{mem}$  so as to fast produce features of the following frames  $[c_{t+1}, \dots, c_{t+n}]$ . The memory network

contains two modules: pyramidal feature attention  $\mathcal{N}_{atten}$  (in Section 3.2) and motion-aided LSTM  $\mathcal{N}_{m-lstm}$  (in Section 3.3.2). Pyramidal feature attention receives  $f_t^l$  as input and generate  $f_t^{atten}$  that will be propagated to the neighboring frames. And the motion-aided LSTM transfers the preceding features by using motion vectors  $m_{t+k}$  and residual errors  $r_{t+k}$ . The above procedure is formulated as:

$$f_t^l = \mathcal{N}_{feat}(I_t) \quad (1)$$

$$c_{t+k} = \begin{cases} \mathcal{N}_{atten}(f_t^3, f_t^4, f_t^5) & k = 0 \\ \mathcal{N}_{m-lstm}(c_{t+k-1}, m_{t+k}, r_{t+k}) & 1 \leq k \leq n \end{cases} \quad (2)$$

$$[b_t, b_{t+1}, \dots, b_{t+n}] = \mathcal{N}_{rfcn}([c_t, c_{t+1}, \dots, c_{t+n}]) \quad (3)$$

where  $[c_t, c_{t+1}, \dots, c_{t+n}]$  denotes the concatenation of the features of one GOP. It means that  $\mathcal{N}_{rfcn}$  will receive the features within the same GOP, and predict their bounding boxes  $[b_t, b_{t+1}, \dots, b_{t+n}]$  simultaneously.

#### 3.2. Pyramidal Feature Attention

Previous methods only propagate high-level feature maps to the neighboring frames (“res5c\_relu” in [54]). In other words, the following P-frames only receive the high-level semantic features from the reference frame. It is not friendly for detecting objects at vastly different scales. [26] utilizes the inherent pyramidal hierarchy of deep CNNs to detect multi-scale objects. However, their predictions are independently made on each level. If we employ the method from still images to videos, we should propagate features and make predictions on several scales at each timestamp. On the contrary, we develop a method that adaptively combines the pyramidal features through attention mechanism within the first memory module (see the first block in Figure 3). The combined pyramidal features are then sent to the motion-aided LSTM (see Section 3.3.2).

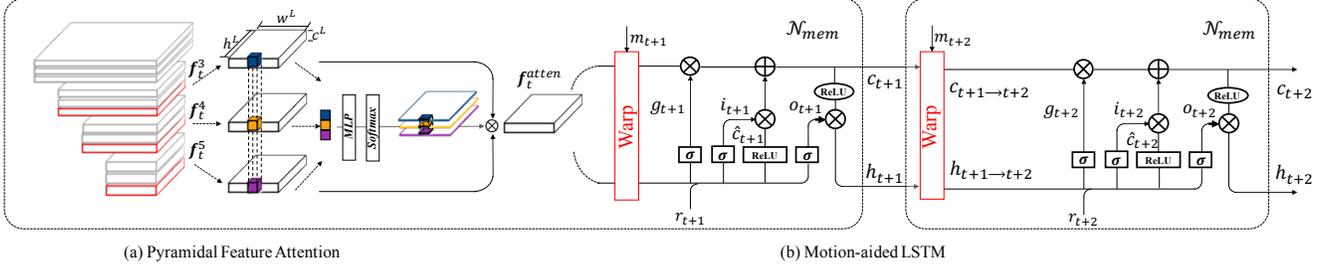


Figure 3. Light-weight MMNet with pyramidal feature attention. Attention mechanism aims to selectively combine the pyramidal features. Motion vectors are used to calibrate cell/hidden features before they run through the memory module. Residual errors are employed to correct appearance changes.

First, we gather the pyramidal features  $\mathbf{f}_t^l$  from different stages. A network can be divided into  $L$  stages whose layers produce output maps of the same resolution. We define one pyramid level for each stage and use the later stages  $l = 3, 4, 5$  of a network (see Figure 3(a)). We utilize  $\{\text{res3b3\_relu}, \text{res4b22\_relu}, \text{res5c\_relu}\}$  of ResNet-101 as the input and transform them into the same dimension:

$$\hat{\mathbf{f}}_t^l = \mathcal{F}_{embed}(\mathbf{f}_t^l, \mathbf{f}_t^L), \forall l, 3 \leq l \leq L \quad (4)$$

where  $\mathcal{F}_{embed}(\cdot)$  can be implemented as a convolution layer with the proper stride. The target dimension is the size of  $\mathbf{f}_t^L$ . Hence,  $\hat{\mathbf{f}}_t^l$  with different  $l$  have the same dimension which is necessary for the following attention operation.

Second, we use squeeze operation across the channel axis to represent features of each scale  $l$  at position  $(i, j)$ :

$$e_t^l(i, j) = \sum_{k=1}^{c^L} \hat{\mathbf{f}}_t^l(i, j, k), \quad (5)$$

$$\forall i, j, 1 \leq i \leq w^L, 1 \leq j \leq h^L$$

where  $i$  and  $j$  enumerate all spatial locations in the feature maps. The squeeze operation sums all the elements across the channel dimension which can be viewed as feature salience. We call the above outputs ‘‘scale descriptors’’. It is inspired by SENet [17], but they use global average pooling to collect statistics for the spatial dimension.

Finally, we adopt the scale descriptors as input to generate attention weights in order to adaptively combine the features from different scales. We define the fused representation  $\mathbf{f}_t^{atten}$  and attention weights  $\alpha_t^l(i, j)$  as follows:

$$\mathbf{f}_t^{atten}(i, j) = \sum_{l=3}^L \alpha_t^l(i, j) \hat{\mathbf{f}}_t^l(i, j) \quad (6)$$

$$\sum_{l=3}^L \alpha_t^l(i, j) = 1$$

The attention weights are produced as follows:

$$\alpha_t(i, j) = \text{softmax}(\text{MLP}(e_t(i, j)))$$

$$e_t(i, j) = [e_t^3(i, j), e_t^4(i, j), e_t^5(i, j)] \quad (7)$$

$$\alpha_t(i, j) = [\alpha_t^3(i, j), \alpha_t^4(i, j), \alpha_t^5(i, j)]$$

After being processed by the pyramidal feature attention,  $\mathbf{f}_t^{atten}$  will be fed into the motion-aligned LSTM.

### 3.3. Motion-aided Memory Network

#### 3.3.1 Motion Vectors and Residual Errors

For compressed data, a P-frame is divided into blocks known as macroblocks (see Figure 4). The supported prediction block sizes range from  $4 \times 4$  to  $16 \times 16$  samples. Video encoder adopts a block matching algorithm [50, 31, 49, 20] to find a block similar to the one it is encoding on a previously encoded frame. The absolute source position and destination position of a macroblock are stored in the motion vectors. Moreover, if there is not an exact match to the block it is encoding, the residual errors are also sent to the decoder.

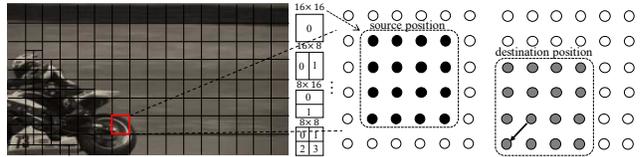


Figure 4. Motion vector represents a macroblock in a picture based on the position of that in another picture.

In our implementation, we use FFmpeg [1] to extract the motion vectors and residual errors for each P-frame. When we obtain the original motion vectors and residuals from the codecs, we resize them to match the size of the feature maps  $h^L$  and  $w^L$ . And the motion vectors should be further rescaled by spatial stride since the original values indicate the movements in the decoded frames.

### 3.3.2 Motion-aided LSTM

We use LSTM [15] to transfer the features. There are two modifications to the conventional LSTMs. One is the motion vector aided feature warping and another is residual error based new input.

Although the gates in LSTM focus on selecting and updating representations, it is still hard for them to forget the object after it has moved to a different position [48]. Experiments in Section 4.4 demonstrate this concern. It is called misaligned features across frames. Hence, we propose a motion vector based feature warping which helps to calibrate cell/hidden features before running through the memory module (see Figure 3(b)). We warp the feature maps from the neighboring frames to the current frame as follows:

$$\begin{aligned} \mathbf{c}_{t+k-1 \rightarrow t+k} &= \mathcal{W}(\mathbf{c}_{t+k-1}, \mathbf{m}_{t+k}) \\ \mathbf{h}_{t+k-1 \rightarrow t+k} &= \mathcal{W}(\mathbf{h}_{t+k-1}, \mathbf{m}_{t+k}) \end{aligned} \quad (8)$$

where  $\mathbf{c}_{t+k-1}$  and  $\mathbf{h}_{t+k-1}$  are outputs of the memory module at time  $t+k-1$ . We set  $\mathbf{c}_t$  and  $\mathbf{h}_t$  to  $\mathbf{f}_t^{\text{atten}}$  and  $k \in [1, n]$ .  $n$  is the number of P-frames in a GOP. The warping operation  $\mathcal{W}$  is similar to [53]. It is implemented by bi-linear function which is applied on each location for all feature maps. It projects a location  $\mathbf{p} + \Delta\mathbf{p}$  in the frame  $t+k-1$  to the location  $\mathbf{p}$  in the frame  $t+k$  which can be formulated as:

$$\begin{aligned} \Delta\mathbf{p} &= \mathbf{m}_{t+k}(\mathbf{p}) \\ \mathbf{c}_{t+k-1 \rightarrow t+k}(\mathbf{p}) &= \sum_q G(\mathbf{q}, \mathbf{p} + \Delta\mathbf{p}) \mathbf{c}_{t+k-1}(\mathbf{q}) \end{aligned} \quad (9)$$

where  $\Delta\mathbf{p}$  is obtained through  $\mathbf{m}_{t+k}$ .  $\mathbf{q}$  enumerates all spatial locations in the feature maps  $\mathbf{c}_{t+k-1}$ , and  $G(\cdot)$  denotes bi-linear interpolation kernel as follow:

$$G(\mathbf{q}, \mathbf{p} + \Delta\mathbf{p}) = \max(0, 1 - \|\mathbf{q} - (\mathbf{p} + \Delta\mathbf{p})\|) \quad (10)$$

Hidden features  $\mathbf{h}_{t+k-1 \rightarrow t+k}$  can also be obtained through above operations. Then  $\mathbf{c}_{t+k-1 \rightarrow t+k}$  and  $\mathbf{h}_{t+k-1 \rightarrow t+k}$  are used as the input from previous time to the current memory module.

For conventional LSTMs, the current complete frame will be used as the new information. In our model, we use the residual errors as new input. Through the motion vector, the previous features can be matched to the current state, but the current representation still lacks some information. So the video encoder computes the residual errors, whose values are known as the prediction error and needed to be transformed and sent to the decoder. After spatial alignment, the residual errors can be used as the complementary information which are more crucial than the whole appearance features of the complete image. In order to better match the residual errors from image-level to the feature-level, we use one convolutional layer to rescale the values.

After obtaining the warped features and new input, the

memory can generate the new cell features as follow:

$$\begin{aligned} \mathbf{g}_{t+k} &= \sigma(\mathbf{W}_g(\mathbf{h}_{t+k-1 \rightarrow t+k}, \mathbf{r}_{t+k})), \\ \mathbf{i}_{t+k} &= \sigma(\mathbf{W}_i(\mathbf{h}_{t+k-1 \rightarrow t+k}, \mathbf{r}_{t+k})), \\ \hat{\mathbf{c}}_{t+k} &= \text{ReLU}(\mathbf{W}_c(\mathbf{h}_{t+k-1 \rightarrow t+k}, \mathbf{r}_{t+k})), \\ \mathbf{c}_{t+k} &= \mathbf{g}_{t+k} \otimes \mathbf{c}_{t+k-1 \rightarrow t+k} + \mathbf{i}_{t+k} \otimes \hat{\mathbf{c}}_{t+k} \end{aligned} \quad (11)$$

where  $\oplus$  and  $\otimes$  are element-wise addition and multiplication, and  $\mathbf{W}_g$ ,  $\mathbf{W}_i$  and  $\mathbf{W}_c$  are learnable weights.  $\mathbf{g}_{t+k}$  can be regarded as a selection mask and  $\hat{\mathbf{c}}_{t+k}$  is new information that holds complementary representation.  $\mathbf{c}_{t+k}$  represents the current frame that will be fed into  $\mathcal{N}_{rfcn}$ . Then the hidden features can be generated:

$$\begin{aligned} \mathbf{o}_{t+k} &= \sigma(\mathbf{W}_o(\mathbf{h}_{t+k-1 \rightarrow t+k}, \mathbf{r}_{t+k})), \\ \mathbf{h}_{t+k} &= \mathbf{o}_{t+k} \otimes \text{ReLU}(\mathbf{c}_{t+k}) \end{aligned} \quad (12)$$

Based on this architecture, we can transform former features to the current state and they will be passed to the next step until encountering another new I-frame. Features of one GOP  $[\mathbf{c}_t, \mathbf{c}_{t+1}, \mathbf{c}_{t+1}, \dots, \mathbf{c}_{t+n}]$  will be sent to the detection network  $\mathcal{N}_{rfcn}$ , producing bounding boxes of objects simultaneously.

## 4. Experiments

### 4.1. Dataset preparation and evaluation metrics

We evaluate the proposed MMNet on the ImageNet [35] object detection from video (VID) dataset. It is split into 3862 training and 555 validation videos. It contains 30 classes labeled with ground truth bounding boxes on all frames. We report the evaluation of previous state-of-the-art models on the validation set and use mean average precision (mAP) as the evaluation metric by following the protocols in [21, 53, 54]. VID releases both original videos and decoded frames. Note that all of the previous state-of-the-art methods use decoded frames as input. It is the first time to detect objects on the original videos on VID.

The 30 object categories in ImageNet VID are a subset of 200 categories in the ImageNet DET dataset. We follow previous approaches and train our model on an intersection of ImageNet VID and DET set.

### 4.2. Training and Evaluation

We perform two phase training: 1) the model is trained on the mixture of DET and VID for 12K iterations, with learning rates of  $2.5 \times 10^{-4}$  and  $2.5 \times 10^{-5}$  in the first 80K and 40K iterations, respectively. We use a batch size of 4 on 4GPUs. 2) the motion-aided memory network is integrated into R-FCN, and trained for another one epoch on VID dataset. In this phase, each GPU holds multiple samples in one GOP. It is already introduced by Section 3.1. The feature extractor ResNet101 model is pre-trained for ImageNet classification as default. In both training and testing, we use single scale images with shorter dimension of

Backbone	ResNet-101					
	(a)	(b)	(c)	(d)	(e)	(f)
MV?			✓	✓	✓	✓
Residual?		✓	✓		✓	✓
LSTM?	✓	✓		✓	✓	✓
Pyramidal Attention?						✓
mAP(%)(fast)	27.7	27.3 ↓ <sub>0.4</sub>	38.5 ↑ <sub>10.8</sub>	43.1 ↑ <sub>15.4</sub>	44.2 ↑ <sub>16.5</sub>	43.7 ↑ <sub>16.0</sub>
mAP(%)(medium)	68.2	68.0 ↓ <sub>0.2</sub>	71.2 ↑ <sub>3.0</sub>	71.5 ↑ <sub>3.3</sub>	72.0 ↑ <sub>3.8</sub>	73.4 ↑ <sub>5.2</sub>
mAP(%)(slow)	82.6	82.2 ↓ <sub>0.4</sub>	83.5 ↑ <sub>0.9</sub>	83.0 ↑ <sub>0.4</sub>	83.6 ↑ <sub>1.0</sub>	84.7 ↑ <sub>2.1</sub>
<b>mAP(%)</b>	<b>66.3</b>	<b>66.1</b> ↓ <sub>0.2</sub>	<b>70.3</b> ↑ <sub>4.0</sub>	<b>71.3</b> ↑ <sub>5.0</sub>	<b>72.1</b> ↑ <sub>5.8</sub>	<b>73.0</b> ↑ <sub>6.7</sub>
Speed(fps)	42.1	41.9	51.3	41.9	41.7	40.5

Table 1. Accuracy of different methods on ImageNet VID validation using ResNet-101 feature extraction networks.

600 pixels. For testing we run the whole recognition network only on I-frame and fast predict the bounding boxes for the rest frames.

### 4.3. Ablation Study

In this section, we conduct an ablation study to prove the effectiveness of the proposed network, including motion vectors, residual errors, LSTM and pyramidal feature attention. We use ResNet-101 to extract I-frame features and adopt different ways to propagate features to the following P-frames. The evaluation protocols follow the previous work [53]. They divide the ground truth objects into three groups according to their movement speed. They use object’ averaged intersection-over-union(IoU) scores with its corresponding instances in the nearby frames as the measurement. The lower motion IoU( $< 0.7$ ) indicates the faster movement. Otherwise, the larger Motion IoU ( $score > 0.9$ ) denotes the object moves slowly. Table 1 also shows the accuracy and runtime speed for the models.

**Method (a) and (b):** Method(a) adopts LSTMs to transform features. It is a conventional solution and we regard it as our baseline. However, without explicit motion cues, LSTMs are unable to automatically align features from previous frames, leading to poor performance (66.3% mAP in Table 1), especially for fast moving objects (27.7% mAP). In method (b), without motion alignment, the residual errors even hurt the result (66.1% mAP).

**Method (c) and (d):** These methods utilize motion vector to warp(/align) the features. Residual errors or LSTMs aim to learn complementary features. we find that {motion vector + residual error} is a practical solution because it has the least computational cost (51.3 fps) with comparable accuracy (70.3%). For the fast moving objects, the result is improved from 27.7% to 38.5 %. It also proves that motion information encoded in the compressed videos is valuable

for modeling differences among frames.

**Method (e) and (f):** These methods are based on motion-aided memory network with/without pyramidal feature attention. Method (e) only propagates high-level feature maps of the top layer and Method (d) delivers pyramidal features to the memory network. We find pyramidal features can further improve the performance with little higher runtime complexity.

To sum up, the motion vector and residual errors are necessary for modeling the motion among consecutive frames. They can speed up the detection procedure. Besides, LSTM is employed to filter out unnecessary information and complement new information. Moreover, propagation of pyramidal features can further improve the detection accuracy. Consequently, these modules are capable of promoting the final feature representations collaboratively.

### 4.4. Visualization

**Visualization of Memory.** We attempt to take a deeper look at intermediate features learned by motion-aided memory network. In Figure 5, there are three typical video snippets. For example, in video #2, the left part consists of decoded frames  $\{I_t, \dots, P_{t+2}, \dots, P_{t+5}, \dots, P_{t+7}\}$ . The car in video moves from left to the middle. We compare the visualization results of mis-aligned and motion-aided memory.

The features in the middle part are learned by LSTM. Although the gates within LSTM are designed for learning the changes among historical information and new input, it is incapable of aligning the spatial features across frames. It cannot capture motion cues only depended on appearance features. The right part presents the motion-aided memory network. Features of  $\{P_{t+2}, P_{t+5}, P_{t+7}\}$  are all based on  $I_t$ . The MMNet receives codecs information as input, aligns and corrects the propagated features. The neurals of

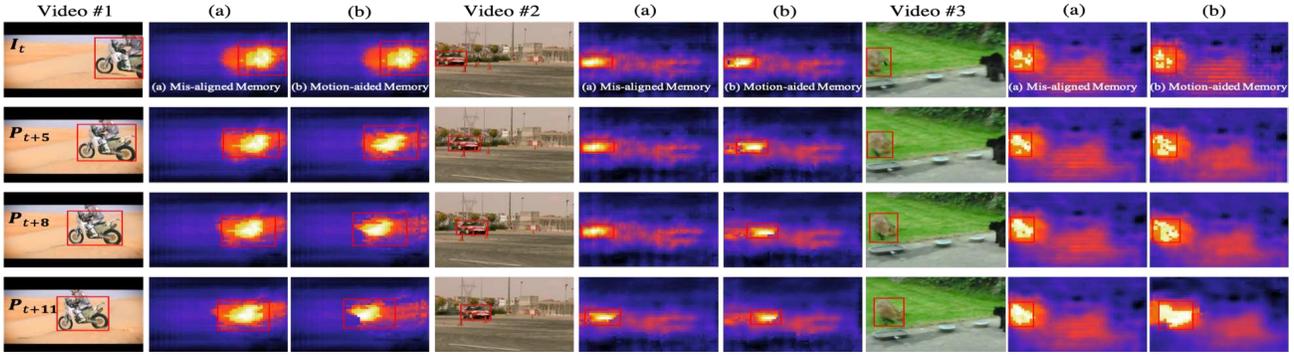


Figure 5. Memory visualization. Each example contains original frames, (a) mis-aligned memory and (b) motion-aided memory. Motion information is quite necessary for feature propagation. It helps MMNet align the feature when the objects move to a different position.

high response in the heatmap move from the left to the middle as the original car.

From the above comparison, motion information is extremely important for feature propagation. It helps to align features when the objects move to a different position. Thus the motion-aided memory network can calibrate features and alleviate inaccurate localizations.

**Visualization of FlowNet and Motion Vector.** In order to show the differences of motion cues between the flow estimation and motion vectors, we visualize two examples and their results in Figure 6. Each of the two examples contains the original short snippet, results of FlowNet [8, 54] and motion vectors (We use the tool provided with Sintel[4]

to visualize the above motion information).

The main advantage of motion vector is freely available. It requires no extra time or models to retrieve motion information because it has already been encoded in the compressed video. From the results in Figure 6, even the motion vector is not as meticulous as FlowNet, it is able to model the motion tendency of objects. All the features of frame  $P_{t+1}$ ,  $P_{t+3}$ ,  $P_{t+5}$ ,  $P_{t+7}$ ,  $P_{t+9}$  are propagated from the I-frame  $I_t$  by utilizing motion vectors, rather than using heavy computational network. Moreover, the bounding boxes location and recognition results are reasonable by the guidance of motion cues, and sometimes even exceed the flow estimation results.



Figure 6. Visualization of FlowNet and Motion Vector. The FlowNet in [54] is capable of building detailed model information. And the motion vector can quickly provide the motion cues which helps to speed up detection procedure in most situations.

Methods	Method	mAP(%)
Single Frame	R-FCN[7]	73.6
Box Propagation	ST-Lattice[5]	77.8
	TCNN[22]	73.8
	Seq-NMS [12]	52.2
	TCN[23]	47.5
Feature Propagation	MANet[44]	78.1
	FGFA[53]	76.5
	DFF[54]	73.1
	TPN[21]	68.4
	Mobile[28]	54.4
	<b>Ours (MMNet)</b>	73.0(41fps) ~76.4(10fps)
	<b>Ours (+PostProc.)</b>	74.8(55fps) ~79.2(8fps)

Table 2. Performance comparison with state-of-the-art systems on the ImageNet VID validation set. The mean average precision (in %) over all classes is shown.

For flow estimation, the motion information is more smooth. It has superior performance when the object is small and unclear. But this model composes dozens of convolutional layers. For each neighboring frame, it should calculate the FlowNet first, which seems not elegant.

To sum up, FlowNet is capable of building detailed motion information. And the motion vector can quickly provide the motion cues which help to speed up detection. This comparison shows the potential of compressed video based detection methods. It fully exploits the codec information and makes the model more elegant.

#### 4.5. Comparison with state-of-the-art systems

In this section, we show the runtime speed and performance of the related methods in Table 2 and Figure 7. In Table 2, the methods are divided into three groups: single frame baseline [7], box-level and feature-level propagation. We also present the detailed accuracy-runtime tradeoff of baseline methods whose performance is above 70% mAP in Figure 7. And the runtime includes the cost of data pre-processing.

From the comparison in Figure 7, we find that:

**Per-frame detection (yellow):** MANet[44] has the best performance among these previous works, whereas it takes about 260ms to detect for one frame. All of these per-frame detectors use heavy computational networks (<10fps);

**Feature propagation (purple):** After producing features on a keyframe, DFF [54] propagates features by using flow estimation. Compared with DFF, our model achieves better performance on both accuracy and runtime speed.

**Box postprocessing (blue):** The box-level propagation is

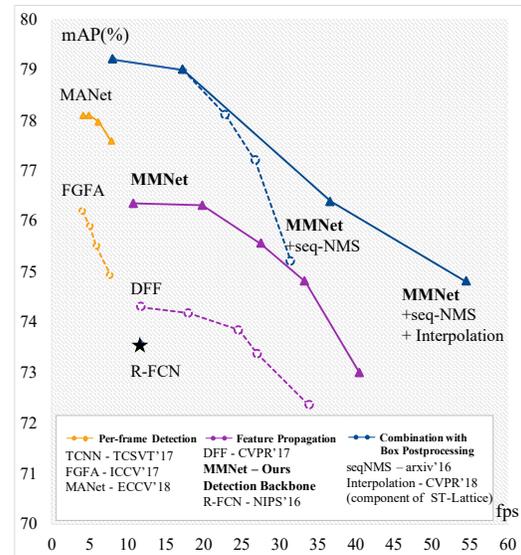


Figure 7. The detailed speed and accuracy of some typical methods. The runtime is measured on an NVIDIA Titan X Pascal GPU.

complementary with feature-level propagation. We select two typical methods seq-NMS[12] and interpolation (part of ST-Lattice) [5] as baselines. When we combine them with our MMNet, they steadily push forward the performance envelope.

**To sum up**, the MMNet performs well in both accuracy and speed, and it can be easily incorporated with box-level post-processing methods.

## 5. Conclusions

In this paper, we propose a fast object detection model incorporating motion-aided memory network called MMNet. It can be directly applied to compressed videos. Different from the previous work, we use motion information stored and transmitted within a video stream, rather than building another model to retrieve motion cues. We use the I-frame as the reference frame, and explore the memory network to transfer features to next P-frames. All these operations are designed with respect to compressed videos. We conduct extensive experiments, including ablation study, visualization and performance comparison, demonstrating the effectiveness of the proposed model.

## Acknowledgments

This work was supported in part by the National Key R&D Program of China under Grant No.2017YFB1302200, by TOYOTA TTAD-2019-08, by DeepBlue AI China, and by Joint Fund of NORINCO Group of China for Advanced Research under Grant No.6141B010318.

## References

- [1] Ffmpeg. <https://github.com/FFmpeg/FFmpeg>.
- [2] Gedas Bertasius, Lorenzo Torresani, and Jianbo Shi. Object detection in video with spatiotemporal sampling networks. In *ECCV*, pages 331–346, 2018.
- [3] Michael Bramberger, Andreas Doblander, Arnold Maier, Bernhard Rinner, and Helmut Schwabach. Distributed embedded smart cameras for surveillance applications. *Computer*, 39(2):68–75, 2006.
- [4] Daniel J Butler, Jonas Wulff, Garrett B Stanley, and Michael J Black. A naturalistic open source movie for optical flow evaluation. In *European Conference on Computer Vision*, pages 611–625. Springer, 2012.
- [5] Kai Chen, Jiaqi Wang, Shuo Yang, Xingcheng Zhang, Yuanjun Xiong, Chen Change Loy, and Dahua Lin. Optimizing video object detection via a scale-time lattice. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7814–7823, 2018.
- [6] Yunpeng Chen, Jianan Li, Huaxin Xiao, Xiaojie Jin, Shuicheng Yan, and Jiashi Feng. Dual path networks. *CoRR*, abs/1707.01629, 2017.
- [7] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. In *Advances in neural information processing systems*, pages 379–387, 2016.
- [8] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2758–2766, 2015.
- [9] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Detect to track and track to detect. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3038–3046, 2017.
- [10] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [11] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.
- [12] Wei Han, Pooya Khorrami, Tom Le Paine, Prajit Ramachandran, Mohammad Babaeizadeh, Honghui Shi, Jianan Li, Shuicheng Yan, and Thomas S Huang. Seq-nms for video object detection. *arXiv preprint arXiv:1602.08465*, 2016.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(9):1904–1916, 2015.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [15] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [16] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [17] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.
- [18] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2261–2269, 2017.
- [19] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016.
- [20] Changsoo Je and Hyung-Min Park. Optimized hierarchical block matching for fast and accurate image registration. *Signal Processing: Image Communication*, 28(7):779–791, 2013.
- [21] Kai Kang, Hongsheng Li, Tong Xiao, Wanli Ouyang, Junjie Yan, Xihui Liu, and Xiaogang Wang. Object detection in videos with tubelet proposal networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 727–735, 2017.
- [22] Kai Kang, Hongsheng Li, Junjie Yan, Xingyu Zeng, Bin Yang, Tong Xiao, Cong Zhang, Zhe Wang, Ruohui Wang, Xiaogang Wang, et al. T-cnn: Tubelets with convolutional neural networks for object detection from videos. *IEEE Transactions on Circuits and Systems for Video Technology*, 2017.
- [23] Kai Kang, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Object detection from video tubelets with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 817–825, 2016.
- [24] Vadim Kantorov and Ivan Laptev. Efficient feature extraction, encoding, and classification for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2593–2600, 2014.
- [25] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [26] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 936–944, 2017.
- [27] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2999–3007, 2017.
- [28] Mason Liu and Menglong Zhu. Mobile video object detection with temporally-aware feature maps. *arXiv preprint arXiv:1711.06368*, 2017.

- [29] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [30] Jawad Nagi, Frederick Ducatelle, Gianni A Di Caro, Dan Cireşan, Ueli Meier, Alessandro Giusti, Farrukh Nagi, Jürgen Schmidhuber, and Luca Maria Gambardella. Max-pooling convolutional neural networks for vision-based hand gesture recognition. In *2011 IEEE International Conference on Signal and Image Processing Applications (ICSIPA)*, pages 342–347. IEEE, 2011.
- [31] Yao Nie and Kai-Kuang Ma. Adaptive rood pattern search for fast block-matching motion estimation. *IEEE Transactions on Image processing*, 11(12):1442–1449, 2002.
- [32] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016.
- [33] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [34] Iain EG Richardson. H264/mpeg-4 part 10 white paper—prediction of intra macroblocks. *Internet Citation, Apr*, 30, 2003.
- [35] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [36] Abhinav Shrivastava, Abhinav Gupta, and Ross B. Girshick. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 761–769, 2016.
- [37] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [38] Anastasis A. Sofokleous. Review: H.264 and MPEG-4 video compression: Video coding for next-generation multimedia. *Comput. J.*, 48(5):563, 2005.
- [39] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A. Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 4278–4284, 2017.
- [40] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [41] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016.
- [42] B. Ugur Töreyn, A. Enis Çetin, Anil Aksay, and M. Bilgay Akhan. Moving object detection in wavelet compressed video. *Sig. Proc.: Image Comm.*, 20(3):255–264, 2005.
- [43] Jasper R. R. Uijlings, Koen E. A. van de Sande, Theo Gevers, and Arnold W. M. Smeulders. Selective search for object recognition. *International Journal of Computer Vision*, 104(2):154–171, 2013.
- [44] Shiyao Wang, Yucong Zhou, Junjie Yan, and Zhidong Deng. Fully motion-aware network for video object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 542–557, 2018.
- [45] Zhenyang Wang, Zhidong Deng, and Shiyao Wang. Accelerating convolutional neural networks with dominant convolutional kernel and knowledge pre-regression. In *European Conference on Computer Vision*, pages 533–548. Springer, 2016.
- [46] Bichen Wu, Forrest Iandola, Peter H Jin, and Kurt Keutzer. Squeezenet: Unified, small, low power fully convolutional neural networks for real-time object detection for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 129–137, 2017.
- [47] Chao-Yuan Wu, Manzil Zaheer, Hexiang Hu, R Manmatha, Alexander J Smola, and Philipp Krähenbühl. Compressed video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6026–6035, 2018.
- [48] Fanyi Xiao and Yong Jae Lee. Video object detection with an aligned spatial-temporal memory. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 485–501, 2018.
- [49] Jiheng Yang, Baocai Yin, Yanfeng Sun, and Nan Zhang. A block-matching based intra frame prediction for h. 264/avc. In *2006 IEEE International Conference on Multimedia and Expo*, pages 705–708. IEEE, 2006.
- [50] Shan Zhu and Kai-Kuang Ma. A new diamond search algorithm for fast block-matching motion estimation. *IEEE transactions on Image Processing*, 9(2):287–290, 2000.
- [51] Xizhou Zhu, Jifeng Dai, Lu Yuan, and Yichen Wei. Towards high performance video object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7210–7218, 2018.
- [52] Xizhou Zhu, Jifeng Dai, Xingchi Zhu, Yichen Wei, and Lu Yuan. Towards high performance video object detection for mobiles. *arXiv preprint arXiv:1804.05830*, 2018.
- [53] Xizhou Zhu, Yujie Wang, Jifeng Dai, Lu Yuan, and Yichen Wei. Flow-guided feature aggregation for video object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 408–417, 2017.
- [54] Xizhou Zhu, Yuwen Xiong, Jifeng Dai, Lu Yuan, and Yichen Wei. Deep feature flow for video recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2349–2358, 2017.
- [55] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8697–8710, 2018.