

# Rate-Accuracy Trade-Off In Video Classification With Deep Convolutional Neural Networks

Mohammad Jubran, Alhabib Abbas, Aaron Chadha and Yiannis Andreopoulos, *Senior Member, IEEE*

**Abstract**—Advanced video classification systems decode video frames to derive the necessary texture and motion representations for ingestion and analysis by spatio-temporal deep convolutional neural networks (CNNs). However, when considering visual Internet-of-Things applications, surveillance systems and semantic crawlers of large video repositories, the video capture and the CNN-based semantic analysis parts do not tend to be co-located. This necessitates the transport of compressed video over networks and incurs significant overhead in bandwidth and energy consumption, thereby significantly undermining the deployment potential of such systems. In this paper, we investigate the trade-off between the encoding bitrate and the achievable accuracy of CNN-based video classification models that directly ingest AVC/H.264 and HEVC encoded videos. Instead of retaining entire compressed video bitstreams and applying complex optical flow calculations prior to CNN processing, we only retain motion vector and select texture information at significantly-reduced bitrates and apply no additional processing prior to CNN ingestion. Based on three CNN architectures and two action recognition datasets, we achieve 11%–94% saving in bitrate with marginal effect on classification accuracy. A model-based selection between multiple CNNs increases these savings further, to the point where, if up to 7% loss of accuracy can be tolerated, video classification can take place with as little as 3 kbps for the transport of the required compressed video information to the system implementing the CNN models.

**Index Terms**—Video classification, convolutional neural networks, video streaming

## I. INTRODUCTION

Action or event recognition and video classification for visual Internet of Things (IoT) systems [1]–[3], video surveillance [4], and fast analysis of large-scale video libraries [5] have been advancing rapidly due to the advent of deep convolutional neural networks (CNNs). Given that such CNNs are very computationally and memory intensive, they are not commonly deployed at the video sensing nodes of the system (a.k.a., “edge” nodes). Instead, video is either transported to certain high-performance aggregator nodes in the network [1]–[3] that carry out the CNN-based processing, or compact features are precomputed in order to allow for less complex

MJ is with the Dept. of Electrical Engineering, Birzeit University, West Bank, Palestine. AA and AC are with the Electronic and Electrical Engineering Department, University College London, Torrington Place, London, WC1E 7JE, UK (e-mail: {alhabib.abbas.13, aaron.chadha.14}@ucl.ac.uk). YA is with the Electronic and Electrical Engineering Department, University College London, Torrington Place, London, WC1E 7JE, UK, and also with iSize Ltd., 41 Corsham Street, London, N1 6DR, UK, [www.isize.co](http://www.isize.co) (e-mail: yiannis@isize.co). We acknowledge support from: the Leverhulme Trust (RAEng/Leverhulme Senior Research Fellowship of YA) and the Royal Commission for the Exhibition of 1851 (Fellowship of AC). MJ performed the work while visiting University College London under a “Distinguished Scholar Award” from the Arab Fund Fellowships programme. This work has been presented in part at the 2018 IEEE Int. Conf. on Image Processing.

on-board processing at the edge [5], usually at the expense of some accuracy loss for the classification or recognition task.

Motion vector based optical flow approximations have been proposed for action recognition by Kantorov and Laptev [6], albeit without the use of CNNs. In more recent work, proposals have been put forward for fast video classification based on CNNs that ingest compressed-domain motion vectors and selective RGB texture information [7], [8]. Despite their significant speed and accuracy improvements, none of these approaches considered the trade-off between rate and classification accuracy obtained from a CNN. Conversely, while rate-accuracy trade-offs have been analysed for conventional image and video spatial feature extraction systems [2], [3], these studies do not cover deep CNNs and semantic video classification, where the different nature of the spatio-temporal classifiers can lead to different rate-accuracy trade-offs.

In this paper, we show that crawling and classification of remote video data can be achieved with significantly-reduced bitrates by exploring rate-accuracy trade-offs in CNN-based classification. Our contributions are summarized as follows:

- 1) We study the effect of varying encoding parameters on state-of-the-art CNN-based video classifiers. Unlike conventional rate-distortion curves, we show that, without any optimization, rate-accuracy is not monotonic for CNN-based classification.
- 2) In order to optimize the trade off between bitrate and classification accuracy, we propose a mechanism to select amongst 2D/3D temporal CNN and spatial CNN classifiers that have varied input volume requirements. We achieve this with minimal modifications to the encoded bitstream, which are straightforward to implement in practice.
- 3) We study and compare the efficacy of our method on action recognition based on AVC/H.264 and HEVC compressed video, which represent two of the most commonly-used video coding standards.

These contributions extend on our recent conference paper on this subject [9], which did not cover the last two points from above. The remainder of the paper is organized as follows. In Section II, we give an overview of recent work on compressed video classification. Section III details how we reduce video bitstreams through selective cropping. In Section IV, we describe and formulate the optimized classifier selection process. Section V evaluates the performance of the proposed classifiers using different coding settings and illustrates the rate gains made possible through our classifier selection method. Finally, Section VI concludes the paper.

## II. RELATED WORK

The use of codec motion vectors as an approximation of optical flow has been proposed for action recognition by Kantorov and Laptev [6]. Their approach preceded the surge in convolutional neural networks for image classification and used Fisher vectors, which achieve lower accuracies in standard action recognition datasets. Recently, Zhang *et al.* [8] utilized codec motion vectors as input to a 2D CNN for low latency video action recognition. They show that distilling knowledge from supervised teachers that operate with optical flow can help student networks learn better representations from compressed video bitstreams. However, their requirement for highly-upsampled motion vector frames during inference [8] increases the implementation complexity, as large activation maps need to be calculated at the first layers of their CNN. Recent work shows that useful features can also be transferred from self-supervised architectures and generative models for visual feature synthesis [10], [11]. Other work [7], [12] showed that high accuracy in compressed-domain action recognition can be achieved without upsampling frames to offer a higher ingestion and CNN processing speed than all previous alternatives. Given that in two-stream architectures [13], [14] the spatial stream learns on scene information that tends to be persistent across frames, compressed-domain methods gain by super-positioning decoded macroblock motion vector information to generate intermediate frames at a finer temporal scale.

However, thus far, there has been no work on exploring rate-accuracy tradeoffs for CNN-based video classification. This is now increasingly important due to the advent of visual IoT and cloud-based platforms, where the visual sensing and processing are not co-located [1]–[3]. Alas, such tradeoffs are non trivial, because they depend on the spatio-temporal information needed by the CNN performing the recognition task [15], [16]. For instance, one of the issues with most of the work described above is the short temporal extent of inputs [7], [17], [18]; each input video segment comprises a small group of frames that only represent approximately one second of the recorded action or event to be classified. Hence, this cannot account for cases where temporal dependencies extend over longer durations [7]. Feichtenhofer *et al.* [19] attempted to resolve this issue by using multiple copies of their two stream network where the copies are spread over a coarse temporal scale, thus encompassing both coarse and fine motion information with an optical flow input. The architecture is spatially and then temporally fused using 3D convolution and pooling. Despite achieving state-of-the-art results on UCF-101 and HMDB-51 datasets, this approach requires heavy processing for both training and testing. Alternatively, other work [15], [20] argues that increasing the temporal extent is simply a case of taking the optical flow component over a larger temporal extent. In order to minimize the complexity of the network, most such approaches downsize the frames, thus reducing the spatial dimensions. On the other hand, the work of Sevilla *et al.* [21] shows that high-resolution optical flow can be beneficial since deep learning methods can learn features from small details. This observation suggests that high-

resolution optical flow can be leveraged to lower the temporal extent of inputs. Understanding the trade-offs in compressed-domain spatio-temporal information and exploring the rate-accuracy characteristics of CNN-based video classification is the objective of this paper.

## III. CROPPED VIDEO BITSTREAMS

We base our reduced-bitstream encoder on the JM reference software of AVC/H.264 [22] and the HM reference software of HEVC [23]. Our modifications to the reference encoders are designed such that the bitrate of the compressed bitstream is kept at a minimum while preserving the information needed to classify videos. Namely, the compressed bitstream should exclusively hold: (i) key texture components corresponding to rapidly-changing input content; (ii) inter-frame predicted macroblocks and their motion compensation parameters; (iii) control signals and headers needed to comply with its corresponding standard.

### A. Summary of Spatio-Temporal Representations in Video Coding Standards

Before applying inter-frame prediction, AVC/H.264 pictures are split into  $16 \times 16$  pixel macroblocks (MB) to represent luminance and chrominance samples, with the chrominance samples further split into  $8 \times 8$  chroma blocks for the widely used 4:2:0 chroma sampling. Macroblocks are the core of the coding layer and form the basis for adaptive inter and intra predictions. Each of the inter-predicted macroblocks is then encoded using blocks from the set  $\{16 \times 16, 16 \times 8, 8 \times 16, 8 \times 8\}$  [23], [24]. The HEVC standard takes on a more adaptive approach and introduces a Coding Tree Unit (CTU) which consists of luma and chroma Coding Tree Blocks (CTB). The size of each luma CTB is drawn from the set  $\{16 \times 16, 32 \times 32, 64 \times 64\}$  where larger size blocks result in better compression efficiency. Iterative partitioning is then applied to divide CTBs into smaller Coding Blocks (CB) resulting in a tree-like structure [25]. The minimum allowed CB size is also specified, this serves as a hyper-parameter to control the granularity of the tree structure produced, this parameter is commonly referred to as *depth* [26].

In both standards, blocks are predicted via translational motion vectors (MVs) that represent the displacement from matching blocks in previous or subsequent reference frames. Increasing the number of small-size blocks increases the granularity of the MV grid at the expense of lower coding efficiency. These MVs represent the temporal activity and have been shown to be highly correlated with optical flow estimates [7]. If the area covered by the MB is static, the MB is “skipped” and is not encoded. The resulting prediction residual from temporal prediction of non-skipped MBs is encoded using transform coding. The transform coefficients are then quantized based on the quantization parameter (QP). The value of the QP per frame can be chosen from 52 values in  $[0, 51]$ , with lower values indicating high-fidelity encoding.

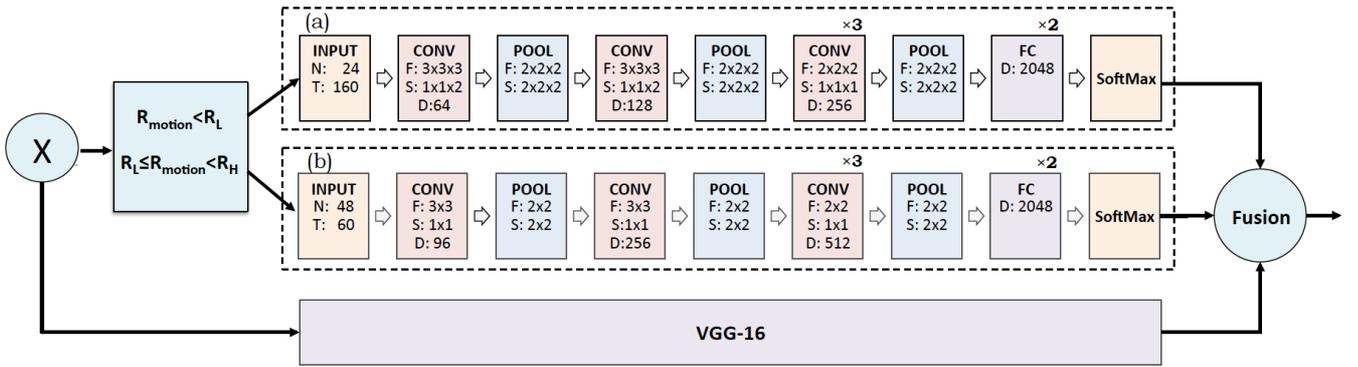


Fig. 1: The proposed Multi-CNN classifier selection: (a) 3D temporal CNN architecture; (b) 2D temporal CNN architecture. The bottom part represents the spatial CNN (VGG-16) [27]. Parameters: N is the spatial dimensions of the input volume; T is the temporal extent expressed as the number of frames used; F is the filter size, formatted as width  $\times$  height  $\times$  time; S is the convolutional window stride; D is the number of filters (or number of hidden units) for the convolutional and fully-connected layers;  $R_L$  and  $R_H$  are controlling the multi-CNN selection based on the motion vector rate  $R_{\text{motion}}$  (Section IV-C).

### B. Selective Retention of Motion and Texture Information

In our work, only select subsets of the quantized transform coefficients are entropy encoded to be included in the cropped bitstream. This set of coefficients, along with spatial texture, is transmitted to the classifier (described in Section IV) to infer semantic features and classify the content of the bitstream. By doing so, the bitrate of these “cropped” subsets of coefficients  $R_{\text{cropped}}$  is significantly reduced in comparison to the original bitrate  $R_{\text{orig}}$  needed to encode the full video. In the remainder of this paper we present our modifications assuming that the first frame of every video sequence is encoded as an Instantaneous Decoding Refresh (IDR), and all subsequent frames in the video are encoded as P-frames.

In order to reduce the bitrate of the compressed bitstream, we employ selective retention of texture information by retaining the texture information of active regions. To implement selective writing in the AVC/H.264 JM reference software [22], we modified functions `writeCoeff4x4_CAVLC_normal()` and `write_chroma_intra_pred_mode()`. In addition, to allow for a skip symbol for all non-active areas, we modified the functions `read_coeff_4x4_CAVLC()` and `read_coeff_4x4_CAVLC_444()`. Similarly, to implement selective writing in the HEVC HM reference software [26], we modified the functions `TEncSbac::codeCoeffNxN()` and `TDecSbac::parseCoeffNxN()`. To simplify our tests, we retain the texture of IDR frames and skip all texture of P-frames with a single skip symbol. The introduction of these skip symbols is the only non-normative part of our entire process. All other syntax elements (including modes and motion information) are left as specified in their respective standard. With these minimal changes, standard decoders can decode our reduced bitstreams to pass to compressed video classifiers.

Finally, in order to derive a temporal activity map from P-frame MVs, we apply the following steps: (i) MVs are extracted from the compressed bitstream using

the `read_motion_info_from_NAL_p_slice()` function for JM and `TDecEntropy::decodePUWise()` for HM; (ii) the extracted MVs are then mapped to a grid of  $8 \times 8$  non-overlapping blocks within each frame; (iii) MVs are interpolated from neighboring macroblocks wherever a macroblock does not provide motion compensation parameters but two or more of its neighbors do.

## IV. PROPOSED FRAMEWORK FOR COMPRESSED-DOMAIN CLASSIFICATION

### A. CNN Architectures

In Fig. 1 we illustrate the two CNNs used for the temporal MV stream, which represent the state-of-the-art in compressed-domain deep learning for action classification [7] [8]. We use two architectures to study how different models behave to cropped bitstream volumes, and to demonstrate that our rate optimized CNN-based classification method is applicable with different network architectures that have been shown to perform well with codec motion vector data. The first CNN architecture we consider is the 3D CNN proposed by Chadha *et al.* [7]. As illustrated in Fig. 1(a), all convolutional and pooling layers are spatiotemporal in extent; this captures the motion information between consecutive motion vector frames. Crucially, the spatiotemporal features are expected to improve classification performance between similar actions. We generate a 4D motion vector input by splitting the  $dx$  and  $dy$  vector components into separate channels, thus resulting in a  $W \times H \times 2 \times T$  volume. We compensate for the low resolution of the extracted motion vector frames by setting a long temporal extent  $T$  as  $T_{3D} = 160$ , which typically comprises the entire video duration.

Similar to [8], the second architecture we consider is a 2D CNN as illustrated in Fig. 1(b). The model design is based on ClarifaiNet [28] and only comprises 2D spatial filters; we notably reduce the size of the first filter from  $7 \times 7$  to  $3 \times 3$  and decrease the stride of the first two convolutional layers to  $1 \times 1$ . The input is generated by stacking the motion vector  $dx$

and  $dy$  components into a single  $W \times H \times 2T$  volume, where the temporal depth  $T$  is set as  $T_{2D} = 60$ . In general, 2D CNNs are less complex to train and test with than 3D CNNs, whilst forgoing modelling any temporal dependencies. Nonetheless, their lower complexity means we can afford to use a higher input spatial resolution, which enables the 2D filters to learn more distinguishing spatial features of the MV data.

Finally, concerning spatial processing of RGB texture, we use the well-established VGG-16 [27] architecture to classify RGB frames and capture motion-invariant spatial features of video content. Our spatial CNN is pre-trained on ImageNet [29] and fine-tuned on the training split of UCF-101. The spatial stream ingests the decoded frames per video and the predictions made by the spatial CNN are ultimately fused with the predictions from the temporal stream to produce the final two-stream classifier decisions.

### B. Training and Testing

We train both temporal stream architectures using stochastic gradient descent with momentum set to 0.9. The initialization of He *et al.* [30] is used and weights are initialized from a normal distribution. Mini-batches of size 64 are generated by randomly selecting 64 training videos per batch. The learning rate is initially set to  $10^{-2}$  and is decreased by a factor of 0.1 every 30k iterations. The training is completed after 90k iterations. We follow the data augmentation practices utilized in recent work [7] in order to minimize overfitting for both the 2D and 3D CNN. These include a multi-scale random cropping of the input and spatial resizing to a fixed size  $N$ , followed by zero centering the motion vector field by subtracting the mean motion vector from the volume. For the 3D CNN, the fixed crop size is set to 24, whereas for the 2D CNN this is doubled to 48. In addition, we use a dropout ratio to 0.5 for the first two fully connected layers in both models. During testing, for the temporal stream we generate 10 random volumes of temporal size  $F$  from which to test on. Per volume, we use the standard 10-crop testing, cropping the four corners and the center of the image to spatial size  $N \times N$  and considering both horizontally flipped and unflipped versions. As such, we average the scores over 10 crops and 10 volumes to produce a single score for the video. For the spatial stream, we use one IDR frame for each video and oversample inputs to VGG-16 by flipping and extracting crops.

### C. Multi-CNN Classifier

In order to optimize the tradeoff between bitrate and classification accuracy, we leverage the differences in input requirements of the two temporal classifiers of Fig. 1 and devise a Multi-CNN (MCNN) selection process. Since the number of MV frames per crop is larger for our 3D CNN versus its 2D CNN counterpart (i.e.,  $T_{3D} > T_{2D}$ ), the former requires higher bitrate per crop than the latter. On the other hand, as shown in previous studies [21], denser MV frames will benefit from the spatially-larger input of the proposed 2D CNN architecture. Since the density of inputs to the temporal stream is directly proportional to the average bitrate allocated to MVs by the codec  $R_{\text{motion}}$ , we expect the accuracies of

both the 2D CNN and 3D CNN classifiers to be correlated to  $R_{\text{motion}}$ , albeit up to a limit (since noise is introduced at high rates due to the limitations of the MV block model). Moreover, the two classifiers are expected to be comparable in accuracy over a range of  $R_{\text{motion}}$  values. These hypotheses have been tested and we present the related experimentally derived results in the Appendix. In summary, our investigation showed that: (i) the long temporal extent 3D CNN classifier is superior for lower values of  $R_{\text{motion}}$ ; (ii) the short temporal extent 2D CNN classifier performs as well as the long temporal extent 3D CNN classifier for mid-range values of  $R_{\text{motion}}$ ; (iii) both temporal CNNs offer diminishing performance for high values of  $R_{\text{motion}}$ . Therefore, we introduce the pair of rate-accuracy optimization parameters  $\{R_L, R_H\}$ , with  $R_H > R_L$ , such that:

- the 3D CNN is used for videos with  $R_{\text{motion}} < R_L$
- the 2D CNN is used for videos with  $R_L \leq R_{\text{motion}} < R_H$
- no temporal CNN is used when  $R_{\text{motion}} \geq R_H$  and only the output of the spatial CNN is considered (see Fig. 1).

The remainder of this section is to establish a model-based approach for the optimal selection of  $\{R_L, R_H\}$ . While the value of  $R_{\text{motion}}$  is derived experimentally during the encoding of each video, for offline rate-accuracy optimization studies it can also be derived via rate-distortion models [31].

### D. Problem Formulation and Optimization of MCNN

To make full use of the overlap of performance between classifiers, a video is passed to a lower-rate classifier only when it is likely to be classified correctly. We consider the problem of finding the optimum set  $\{R_L^*, R_H^*\}$  that maximizes the classification accuracy,  $A_{\text{mcnn}}$ , of our proposed MCNN under a constraint on the available bitrate,  $R_{\text{available}}$ :

$$\{R_L^*, R_H^*\} = \arg \max_{R_L, R_H} A_{\text{mcnn}} \quad \text{subject to} \quad R_{\text{sent}} \leq R_{\text{available}} \quad (1)$$

where  $R_{\text{sent}}$  is the average bitrate of all transmitted bitstreams under a selection algorithm for  $\{R_L, R_H\}$ . We first consider the video source probability density function  $f_s(R_{\text{motion}})$ , which characterizes the probability of occurrence of video examples with bitrate  $R_{\text{motion}}$ . We have found  $f_s(R_{\text{motion}})$  to be well approximated by the Gamma distribution,  $f_s(R_{\text{motion}}; \alpha, \beta)$ , where  $\alpha$  and  $\beta$  are the shape and rate parameters (see Section A of the Appendix and Fig. 7). We can then express  $A_{\text{mcnn}}$  as:

$$\begin{aligned} A_{\text{mcnn}} = & A_{3D} \int_0^{R_L} f_s(R_{\text{motion}}) dR_{\text{motion}} \\ & + A_{2D} \int_{R_L}^{R_H} f_s(R_{\text{motion}}) dR_{\text{motion}} \\ & + A_{\text{SP}} \int_{R_H}^{\infty} f_s(R_{\text{motion}}) dR_{\text{motion}} \end{aligned} \quad (2)$$

where  $A_{3D}$ ,  $A_{2D}$  and  $A_{\text{SP}}$  are the classification accuracies of the 3D, 2D and spatial stream classifiers respectively. In (2), the accuracy of each classifier is assumed to be constant for the range of bitrates it corresponds to, and is experimentally derived from the training dataset. This assumption holds as long as the training dataset is large enough and the accuracy of each classifier remains relatively flat within the respective integration interval of each classifier, which is found to be the case in our experiments of Section V.

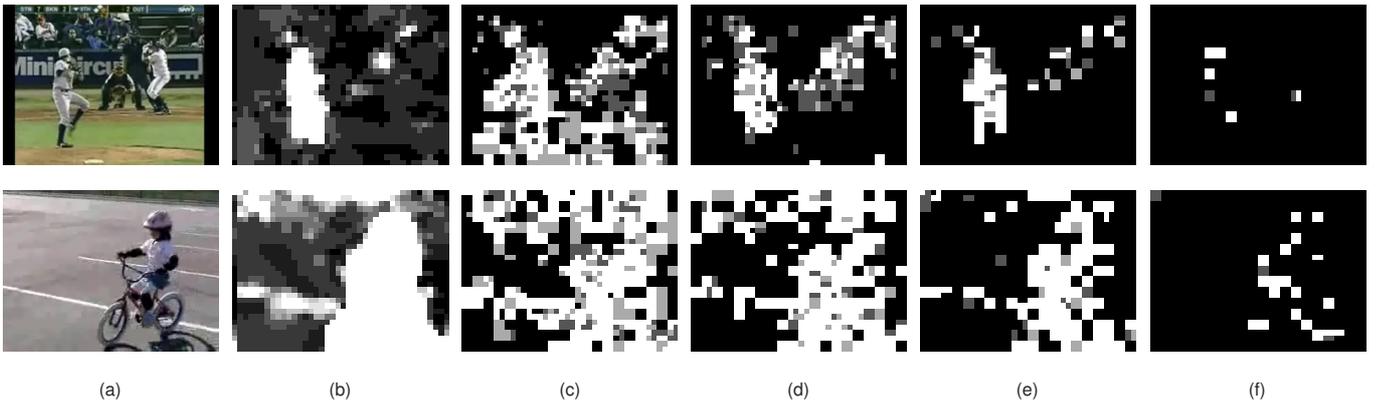


Fig. 2: RGB frames and corresponding AVC/H.264 MV activity maps for two scenes from UCF-101; (a) RGB frames; (b) Brox optical flow; (c) Approximated flow at QP = 0; (d) Approximated flow at QP = 30; (e) Approximated flow at QP = 40; (f) Approximated flow at QP = 51. Note that sparsity increases and noise decreases with increased QP.

Since the number of bits needed to classify each video depends on which classifier is used for prediction, we first find the average bitrate required by each classifier. We define  $R_{3D}$ ,  $R_{2D}$ , and  $R_{SP}$  as the average bitrate of inputs to the 3D, 2D, and spatial classifiers, respectively, and estimate each as:

$$R = \begin{cases} R_{3D} = a_{3D}R_{\text{motion}} + b_{3D} & 0 < R_{\text{motion}} < R_L \\ R_{2D} = a_{2D}R_{\text{motion}} + b_{2D} & R_L \leq R_{\text{motion}} < R_H \\ R_{SP} = I_{SP} & R_H \leq R_{\text{motion}} < \infty \end{cases} \quad (3)$$

where  $a_{3D}$ ,  $b_{3D}$ ,  $a_{2D}$ , and  $b_{2D}$  are coefficients to be estimated by applying regression on the bitrate feature  $R_{\text{motion}}$  obtained on the training set  $\mathcal{V}$ . Since the inputs passed to the 3D and 2D classifiers consist only of the motion vectors and some added headers to comply with the used standard, we expect the linear relations shown in (3) and confirm this in Section B of the Appendix. For the spatial classifier, we use  $I_{SP}$ , i.e., the bitrate of the first IDR frame, to estimate  $R_{SP}$ . Note that  $R_{\text{motion}}$  is not used for  $R_{SP}$ , since the spatial classifier only uses texture information. We can now express  $R_{\text{sent}}$  as:

$$R_{\text{sent}} = \int_0^{R_L} R_{3D} f_s(R_{\text{motion}}) dR_{\text{motion}} + \int_{R_L}^{R_H} R_{2D} f_s(R_{\text{motion}}) dR_{\text{motion}} + \int_{R_H}^{\infty} R_{SP} f_s(R_{\text{motion}}) dR_{\text{motion}} \quad (4)$$

Based on the expectation value property of the Gamma density function  $f(X; \alpha, \beta)$  [32]:

$$Xf(X; \alpha, \beta) = \frac{\alpha}{\beta} f(X; \alpha + 1, \beta) \quad (5)$$

from (2) and (4) we can rewrite  $A_{\text{mcdn}}$  and  $R_{\text{sent}}$  as:

$$A_{\text{mcdn}} = (A_{3D} - A_{2D})F_s(R_L; \alpha, \beta) + (A_{2D} - A_{SP})F_s(R_H; \alpha, \beta) + A_{SP} \quad (6)$$

$$R_{\text{sent}} = (b_{3D} - b_{2D})F_s(R_L; \alpha, \beta) + (b_{2D} - I_{SP})F_s(R_H; \alpha, \beta) + (\alpha/\beta)(a_{3D} - a_{2D})F_s(R_L; \alpha + 1, \beta) + (\alpha/\beta)(a_{2D})F_s(R_H; \alpha + 1, \beta) + I_{SP} \quad (7)$$

where  $F_s$  is the cumulative distribution function of  $f_s$  and we have explicitly indicated the dependence on the parameters  $\alpha$  and  $\beta$  since they affect the bitrate and accuracy contributions of the 2D and 3D CNN models. The constrained optimization problem of (1) can now be solved for  $\{R_L^*, R_H^*\}$  via (6) and (7). We first note that (6) is monotonically increasing in function of  $R_L$  and  $R_H$ , since  $A_{3D} > A_{2D}$  and  $A_{2D} > A_{SP}$ . This allows for the use numerical methods that gradually explore the parameter space of  $\{R_L, R_H\}$  by setting  $R_{\text{sent}}$  in (7) as close as possible to  $R_{\text{available}}$  and then finding the maximum values for  $\{R_L, R_H\}$  that satisfy (7), since such values will automatically maximize (6).

In our experiments, amongst several alternatives, we opted for the method of Toint *et al.* [33], which finds the solution  $\{R_L^*, R_H^*\}$  that maximizes (6) under the constraint  $R_{\text{sent}} \leq R_{\text{available}}$  with the provision of sufficient exploration time. Given that this optimization process is done offline based on training data  $\mathcal{V}$ , this does not impose any overhead at runtime. Finally, we remark that, in case  $R_{\text{motion}}$  is not measurable at training or test time, the optimization method proposed in this section can be generalized to other features that correlate with  $R_{\text{motion}}$  (e.g. number of MVs per frame).

## V. EXPERIMENTAL RESULTS

### A. Used Datasets and Rate Saving from Cropped Bitstreams

We train and test our 2D and 3D CNN architectures on eight distinct motion vector datasets generated by varying the QP setting of AVC/H.264 and HEVC to encode UCF-101 [34], while skipping texture information as described in Section III. For all videos: the first frame is encoded as an IDR (with remaining frames inter-predicted as P-frames), the frame rate is set to 25, and we set the motion vector search range to 16 pixels. Since specifying a particular quantization parameter has a direct effect on the MVs produced by AVC/H.264 and HEVC, this gives several distinct source distributions for the classifier to be trained and tested on.

TABLE I: Average AVC/H.264 bitrate (kbps) of UCF-101;  $R_{\text{orig}}$  is the bitrate of the original bitstream,  $R_{\text{cropped}}$  is the bitrate after cropping and retaining texture and motion information, and  $R_{\text{motion}}$  is the MV bitrate.

QP	$R_{\text{orig}}$	$R_{\text{cropped}}$	$R_{\text{motion}}$	% of $R_{\text{motion}}$ to	
				$R_{\text{orig}}$	$R_{\text{cropped}}$
0	4273.0	321.3	155.4	3.6	48.3
30	274.9	112.3	46.9	17.0	41.7
40	80.0	49.9	18.5	23.2	37.1
51	27.7	20.0	4.6	16.7	23.1

TABLE II: Average HEVC bitrate (kbps) of UCF-101;  $R_{\text{orig}}$  is the bitrate of the original bitstream,  $R_{\text{cropped}}$  is the bitrate after cropping and retaining texture and motion information, and  $R_{\text{motion}}$  is the MV bitrate.

QP	$R_{\text{orig}}$	$R_{\text{cropped}}$	$R_{\text{motion}}$	% of $R_{\text{motion}}$ to	
				$R_{\text{orig}}$	$R_{\text{cropped}}$
0	3065.2	204.9	39.9	1.3	19.1
30	157.7	58.8	12.0	7.6	20.6
40	40.2	26.7	4.9	2.5	12.25
51	10.9	9.8	0.8	7.3	8.1

### B. Rate-Accuracy Results

As the quality of predictions made by CNN models is strongly tied to the properties of the source distribution (e.g. cross-class variance, noise), we expect that varying the rate should affect the accuracy of our classifier accordingly. Since the QP values control the video rate, we first show visual examples of the effect of QP on the quality of approximated sparse optical flow in Fig. 2. The best approximations appear to be for QP values in the region of 30 to 40. To assess the rate savings and classification accuracy of our proposal when varying QP values, in Table I and Table II we compare the original bitrate,  $R_{\text{orig}}$ , with the bitrate of the cropped bitstreams,  $R_{\text{cropped}}$ , and the rate of retained motion vectors,  $R_{\text{motion}}$ . The results show that streaming cropped bitstreams allows for 28% to 92% reduction in bitrate for AVC/H.264, and 11% to 94% for HEVC. The related classification accuracy results are presented in Fig. 3 and Fig. 4. As indicated by the visual examples of Fig. 2, the utilized CNNs indeed achieve their best accuracies at QP values of 30 to 40.

Importantly, we observe that rate-accuracy curves are not monotonic (i.e., accuracy decreases for very low or very high QP values). We expect sparser motion vectors to make certain classes with high motion similarity particularly harder to classify and easier to confuse with each other (e.g., MVs produced by setting QP = 51 where the rate allocated to motion vectors is the lowest). On the other hand, as shown in Fig. 2, setting QP < 30 also has a detrimental effect on accuracy. This is because the derived MVs become significantly more noisy due to the inadequacy of the simple translational block model of AVC/H.264 and HEVC to smoothly approximate the optical flow field, since such block models are optimized for rate control and not for optical flow estimation [7], [35].

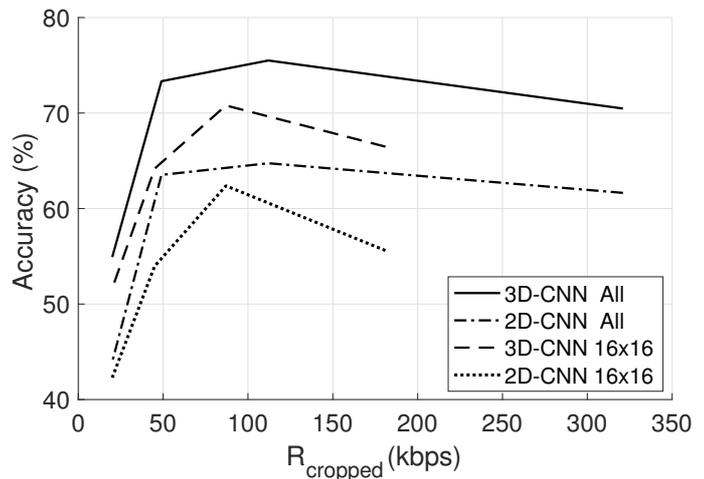


Fig. 3: Rate-accuracy after cropped AVC/H.264 bitstreams are passed to the 2D and 3D classifiers. Each point for every curve corresponds to a different QP setting during encoding, with “16 × 16” indicating restriction to 16 × 16 blocks (no MB subblocks) and “All” indicating the use of all MB partitions.

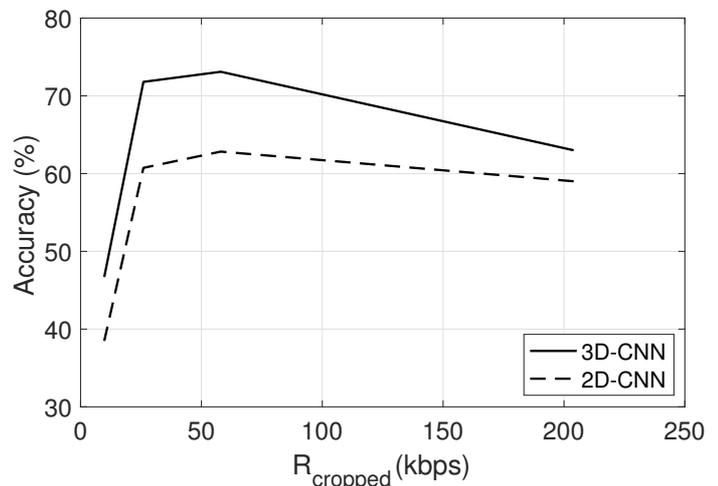


Fig. 4: Rate-accuracy after cropped HEVC bitstreams are passed to the 2D and 3D temporal CNNs. Each point for every curve corresponds to a different QP setting during encoding, with encoder parameter CBT Depth = 2.

To cross validate with an external benchmark, Fig. 5 shows the average End Point Error (aEPE) between MV frames and a dense optical flow ground truth approximated using the method proposed by Brox *et al.* [39]. The resulting curves show that, for both video coders, the minimum aEPE value against dense optical flow is in the QP range of 30 to 40. We also note that the best performance occurs at a lower rate for HEVC compared to AVC/H.264, which is due to the enhanced coding efficiency and improved inter-frame macroblock search of the HEVC standard. This is also reflected in Fig. 5, where the aEPE of HEVC is lower than that of AVC/H.264 over all QP settings.

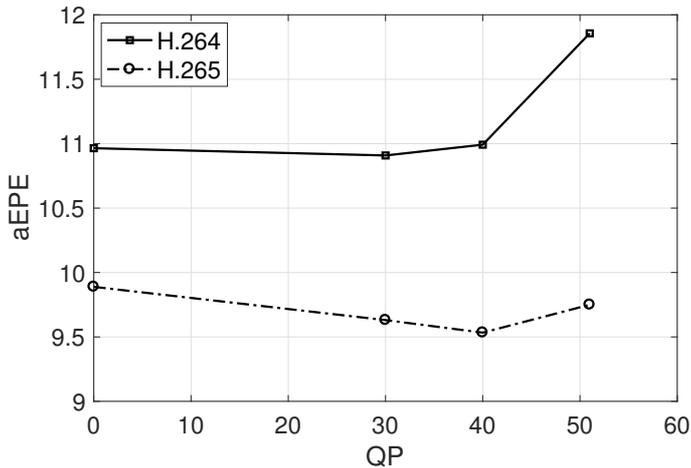


Fig. 5: Average EPE between our approximated optical flow with different QP settings and an estimated dense optical flow ground truth using the method of Brox *et al.* [35].

Framework	$R_{\text{sent}}$ (kbps)	Accuracy (%)	
		UCF	HMDB
MCNN-AVC/H.264, lowest bitrate	2.8	80.4	43.4
MCNN-AVC/H.264, highest bitrate	49.9	88.1	52.9
MCNN-HEVC, lowest bitrate	2.1	80.0	41.8
MCNN-HEVC, highest bitrate	26.7	86.6	50.7
EMV + RGB-CNN [8]	—	86.4	—
MVCNN [7]	—	89.8	56.0
CoViAR [12]	—	90.4	59.1
TCLSTA [36]	—	94.0	68.7
ST-ResNet + iDT [19]	—	94.6	70.3
ActionVLAD + iDT [19]	—	93.6	69.8
TSN (3 modalities) [37]	—	94.2	69.4
I3D [14]	—	93.4	66.4
TSCNN (SVM fusion) [13]	—	88.0	59.4
LTC [20]	—	91.7	64.8
C3D (3 nets)+iDT [38]	—	90.4	—

TABLE III: Comparison of our MCNN classifier (when  $\{R_L, R_H\}$  is selected such that the lowest and highest possible bitrates are achieved) against state-of-the-art CNNs.  $R_{\text{sent}}$  is estimated from UCF-101 and is not relevant for other methods as they require full decompression of whole bitstreams.

### C. MCNN Performance

To study the performance of our proposed MCNN classifier under varying bitrate constraints, as described in Section IV-D, we solve (1) for multiple values of  $R_{\text{sent}}$  within the interval  $[0, 50]$  kbps to assess the accuracy of MCNN on UCF-101 for each set of parameters  $\{R_L^*, R_H^*\}$ . The results are shown in Fig. 6. The lowest bitrate is achieved when  $\{R_L, R_H\}$  is selected such that only the spatial stream is used for all inputs. The highest accuracy is achieved by selecting  $\{R_L, R_H\}$  such that only the two-stream 3D-CNN-F classifier is used. When using the optimization framework of Section IV-D, approximately 25 kbps (50%) reduction in bitrate can be obtained against the 3D-CNN-F classifier (25 kbps vs. 50 kbps) at less than 2% reduction in classification accuracy.

Importantly, further bitrate reductions are made possible with graceful (and monotonic) degradation in classification accuracy, to the point of making it viable to get an accuracy within 7% from the top performance at an average bitrate as low as  $R_{\text{sent}} = 3$  kbps. This shows the potential for further exploration of rate-accuracy optimization in CNN-based video classification and the utility of features such as  $R_{\text{motion}}$  in inferring the amount of information needed for classification.

### D. Comparison Against External Benchmarks

In Table III, we report the accuracy of our MCNN classifier of Fig. 1 when bitstreams with the lowest and highest possible bitrates are allowed, and compare against state-of-the-art methods. Our results show that our approach remains competitive to the state-of-the-art on UCF-101, while retaining the significant bitrate gains reported in Table I and Table II. In addition, while our approach is outperformed by methods like ST-ResNet and TSN, it is important to emphasize that these methods are orders-of-magnitude more complex than operating with sparse compressed-domain information [7], [8], [15], since they require the use of dense optical flow and need to receive and decode entire video bitstreams. Moreover, ST-ResNet and TSN use significantly deeper neural network architectures in comparison to our approach, which makes their inference significantly more compute intensive than the CNN architectures of Fig. 1. For HMDB, MCNN is significantly outperformed by state-of-the-art methods. For example, ST-ResNet outperforms the highest bitrate setting of MCNN-AVC/H.264 by 17.4%. Given that optical flow frames approximated from motion vectors are sparse, and coupled with the motion vector noise naturally occurring in HMDB due to its non-static backgrounds, this induces higher drop in accuracy compared to UCF. This effect is also evidenced in the other motion-vector based methods, i.e., MVCNN [7] and CoViAR [12]. In order to improve our results for HMDB, our rate optimization method could be applied in conjunction with more advanced optical flow estimation methods derived from codec motion vector data (e.g., the recent motion vector accumulation method proposed in CoViAR [12]), albeit at the expense of increased complexity.

## VI. CONCLUSION

We present the first exploration of rate-accuracy trade-offs in advanced video classification with CNNs. Given that our proposed method can be applied based on standardized codecs with minimal bitstream modifications, it is well suited for visual IoT or semantic video crawling applications. The obtained results show that, when reducing bitstreams to the necessary elements for 2D or 3D CNN classification, 28%-92% and 11%-94% reduction in bitrate can be achieved for AVC/H.264 and HEVC respectively. We have observed that non-monotonic rate-accuracy curves are obtained by state-of-the-art CNNs when classifying approximated flow from compressed bitstreams (following the AVC/H.264 and HEVC standards). On the other hand, a rate-based selection method

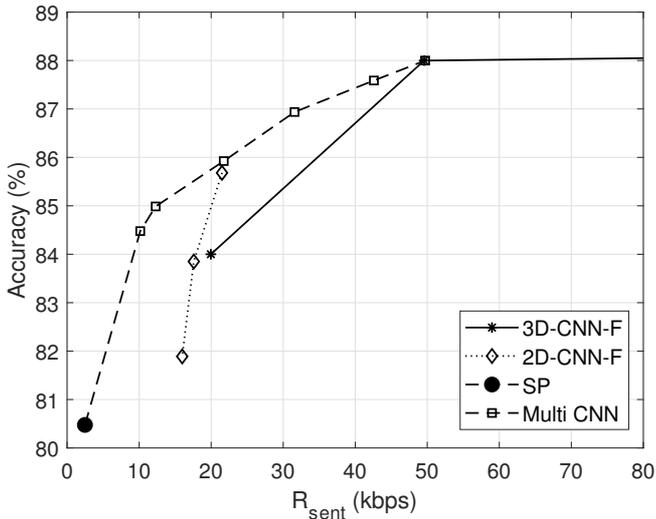


Fig. 6: Rate-accuracy results on the UCF-101 dataset. For the 3D-CNN-F and 2D-CNN-F classifiers (fusion of spatial CNN with 3D/2D motion CNNs as shown in Fig. 1), different rates are obtained by using different QP settings. When using Multi-CNN, rate is controlled by setting QP = 40 and varying  $R_{\text{available}}$  to solve for  $R_L^*$  and  $R_H^*$ . Note that the leftmost point shows the performance when the temporal stream is not used and the MCNN selector only considers the outputs of the spatial stream model.

between multiple CNN classifiers with varied input requirements is shown to achieve monotonic rate-accuracy characteristics and allow for even further rate gains, with minimal impact on classification accuracy. Our implementation and all tools needed to reproduce our results are available online at: <https://github.com/rate-accuracy-mvcnn/main>.

#### APPENDIX

We validate our modelling choices described in Section IV-D. For brevity of exposition, all figures and results here are reported for the indicative case of AVC/H.264 with QP = 40.

##### A. Distribution of $R_{\text{motion}}$ and Performance Overlap

In this section we compare the distribution of  $R_{\text{motion}}$  against the fitted model and verify the overlap of performance between the proposed architectures in Section IV. All of the UCF-101 dataset is used to produce the results shown in Fig. 7 and Fig. 8. For Fig. 7, the Kullback-Leibler divergence (describing the distance between the empirical and fitted Gamma distribution) was found to be 0.034. This proximity justifies our use of this distribution for characterizing the probability of occurrence of different values of  $R_{\text{motion}}$ . Concerning Fig. 8, the experiments show that the 3D and 2D CNN architectures perform similarly for middle-range values of  $R_{\text{motion}}$ , with the 3D-CNN outperforming the 2D-CNN for most of the lower MV bitrates. The performance of both CNNs decays for high values of  $R_{\text{motion}}$ . Hence, for the high-end range of  $R_{\text{motion}}$ , only the spatial CNN should be used (VGG-16 of Fig. 1).

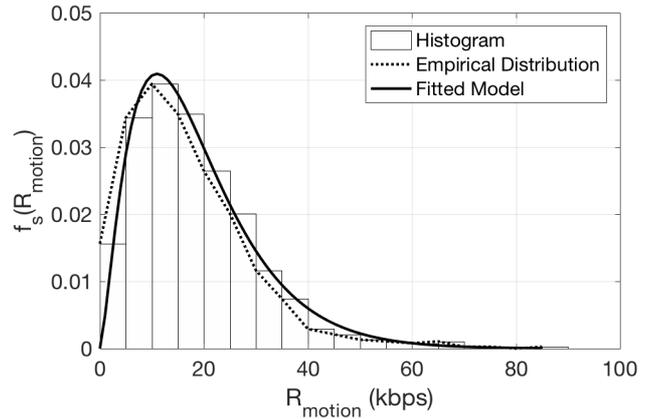


Fig. 7: Empirically measured distribution of  $R_{\text{motion}}$  and fitted Gamma distribution with shape and scale parameters:  $\alpha = 2.43$ ,  $\beta = 0.13$ .

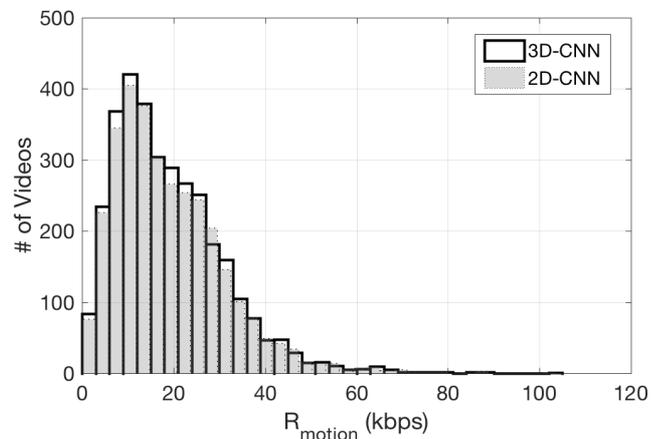


Fig. 8: Number of videos classified correctly by each temporal CNN classifier for different values of  $R_{\text{motion}}$ .

##### B. Linear Model Verification for (3)

We selected 5% of the UCF-101 videos randomly and present the plots of  $R_{\text{motion}}$  vs.  $R_{3D}$  and  $R_{2D}$  in Fig 9 and Fig. 10. Using the same set, we calculated the coefficient of determination  $R^2$  to relate the experimental variance to the residual variance of the linear model and found it to be 93% for  $R_{3D}$  and 88% for  $R_{2D}$ . Similar results have been obtained for the HMDB dataset. These results validate that the linear assumption of (3) is a good approximation.

#### REFERENCES

- [1] J. Posada, C. Toro, I. Barandiaran, D. Oyarzun, D. Stricker, R. de Amicis, E. B. Pinto, P. Eisert, J. Döllner, and I. Vallarino, "Visual computing as a key enabling technology for industrie 4.0 and industrial internet," *IEEE Trans. Computer Graphics and Applications*, vol. 35, no. 2, pp. 26–40, Mar 2015.
- [2] A. Redondi, M. Cesana, and M. Tagliasacchi, "Rate-accuracy optimization in visual wireless sensor networks," *Proc. IEEE Int. Conf. Image Processing (ICIP)*, pp. 1105–1108, 2012.
- [3] A. Redondi, L. Baroffio, M. Cesana, and M. Tagliasacchi, "Compress-then-analyze vs. analyze-then-compress: Two paradigms for image analysis in visual sensor networks," *presented at the 15th Int. Workshop on Multimedia Signal Processing (MMSP)*, pp. 278–282, 2013.

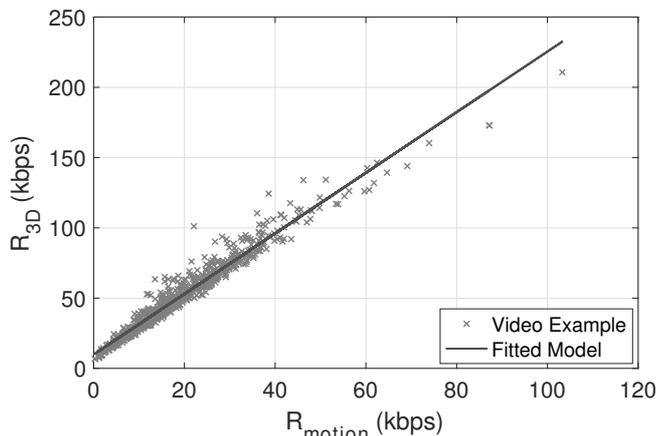


Fig. 9: Bitrate of inputs sent to 3D architecture  $R_{3D}$  plotted against  $R_{\text{motion}}$  and fitted model of  $R_{3D}$  with linear coefficients  $a_{3D} = 2.21$  and  $b_{3D} = 9.04$ .

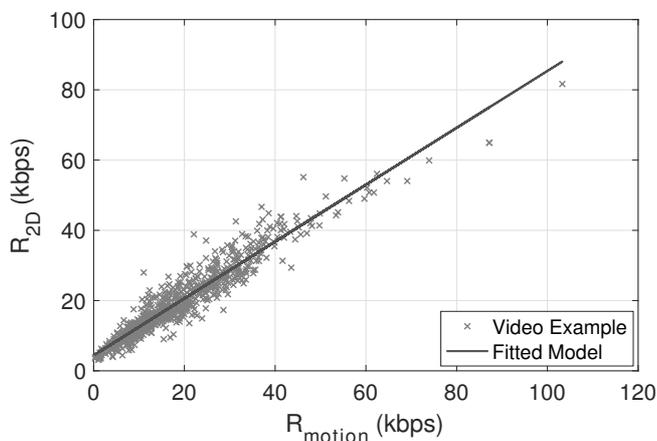


Fig. 10: Bitrate of inputs sent to 2D architecture  $R_{2D}$  plotted against  $R_{\text{motion}}$  and fitted model of  $R_{2D}$  with linear coefficients  $a_{2D} = 0.83$  and  $b_{2D} = 4.27$ .

- [4] G. Yu and J. Yuan, "Fast action proposals for human action detection and search," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pp. 1302–1311, 2015.
- [5] S. Abu-El-Hajja, N. Kothari, J. Lee, A. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan, "YouTube-8M: A large-scale video classification benchmark," *presented at the Large-Scale Video Understanding Workshop session in conjunction with IEEE Conf. on Vision and Pattern Recognition (CVPR)*, 2017.
- [6] V. Kantorov and I. Laptev, "Efficient feature extraction, encoding and classification for action recognition," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pp. 2593–2600, 2014.
- [7] A. Chadha, A. Abbas, and Y. Andreopoulos, "Video classification with CNNs: Using the codec as a spatio-temporal activity sensor," *IEEE Trans. Circuits and Syst. for Video Technol.* DOI:10.1109/TCSVT.2017.2786999, to appear.
- [8] B. Zhang, L. Wang, Z. Wang, Y. Qiao, and H. Wang, "Real-time action recognition with deeply transferred motion vector CNNs," *IEEE Trans. on Image Processing*, vol. 27, no. 5, pp. 2326–2339, May 2018.
- [9] A. Abbas, M. Jubran, A. Chadha, and Y. Andreopoulos, "Rate-accuracy trade-off in video classification with deep convolutional neural networks," *Proc. IEEE Int. Conf. on Image Processing (ICIP)*, pp. 793–797, 2018.
- [10] C. Zhang and Y. Peng, "Better and faster: Knowledge transfer from multiple self-supervised learning tasks via graph distillation for video classification," *Proc. 27th Int. Joint Conf. on Artificial Intelligence (IJCAI)*, pp. 1135–1141, 2018.
- [11] C. Zhang and Y. Peng, "Visual data synthesis via GAN for zero-shot video classification," *Proc. 27th Int. Joint Conf. on Artificial Intelligence (IJCAI)*, pp. 1128–1134, 2018.
- [12] C. Wu, M. Zaheer, H. Hu, R. Manmatha, A. J. Smola, and P. Krähenbühl, "Compressed video action recognition," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pp. 6026–6035, 2018.
- [13] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 568–576, 2014.
- [14] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pp. 4724–4733, 2017.
- [15] R. Girdhar, D. Ramanan, A. Gupta, J. Sivic, and B. Russell, "Action-VLAD: Learning spatio-temporal aggregation for action classification," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pp. 3165–3174, 2017.
- [16] R. Girdhar and D. Ramanan, "Attentional pooling for action recognition," *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 33–44, 2017.
- [17] K. Xu, X. Jiang, and T. Sun, "Two-stream dictionary learning architecture for action recognition," *IEEE Trans. on Circ. and Syst. for Video Technol.*, vol. 27, no. 3, pp. 567–576, Feb 2017.
- [18] S. Zhao, Y. Liu, Y. Han, R. Hong, Q. Hu, and Q. Tian, "Pooling the convolutional layers in deep convnets for video action recognition," *IEEE Trans. on Circ. and Syst. for Video Technol.*, vol. 28, pp. 1839–1849, Mar 2017.
- [19] C. Feichtenhofer, A. Pinz, and A. Zisserman, "Convolutional two-stream network fusion for video action recognition," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pp. 1933–1941, 2016.
- [20] G. Varol, I. Laptev, and C. Schmid, "Long-term temporal convolutions for action recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 40, pp. 1510–1517, 2017.
- [21] L. Sevilla-Lara, Y. Liao, F. Guzey, V. Jampani, A. Geiger, and M. J. Black, "On the integration of optical flow and action recognition," *presented at the German Conference on Pattern Recognition (GCPR)*, Oct 2018.
- [22] T. Alexis, S. Gary, S. Karsten, O. Tobias, and L. Athanasios, "H.264 reference software," <http://iphome.hhi.de/suehring/ttml/>, 2009.
- [23] "Advanced video coding for generic audio-visual services, ITU-T Rec. H.264 and ISO/IEC 14496-10 (AVC)," ITU-T and ISO/IEC JTC 1, 2003.
- [24] T. Wiegand, G. J. Sullivan, G. Bjontegaard, and A. Luthra, "Overview of the H.264/AVC video coding standard," *IEEE Trans. Circ. and Syst. for Video Technol.*, vol. 13, no. 7, pp. 560–576, Jul 2003.
- [25] H. Samet, "The quadtree and related hierarchical data structures," *ACM Comput. Surv.*, DOI:10.1145/356924.356930, vol. 16, no. 2, pp. 187–260, Jun 1984.
- [26] W. J. H. G. J. Sullivan, J. R. Ohm and T. Wiegand, "Overview of the high efficiency video coding (HEVC) standard," *IEEE Trans. Circ. and Syst. for Vid. Technol.*, vol. 22, no. 12, pp. 1649–1668, Dec 2012.
- [27] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *Int. Conf. on Learning Representations (ICLR), oral presentation paper no. 4*, 2015.
- [28] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [29] J. Deng *et al.*, "ImageNet: A large-scale hierarchical image database," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pp. 248–255, 2009.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification," *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, pp. 1026–1034, 2015.
- [31] K. Kwon, Y. Shen, and J. Kuo, "Rate control for H.264 video with enhanced rate and distortion models," *IEEE Trans. Circuits and Syst. for Video Technol.*, vol. 17, no. 5, pp. 517–529, May 2007.
- [32] M. H. DeGroot and M. J. Schervish, *Probability and statistics*. Pearson Education, 2012.
- [33] A. Conn, N. Gould, and P. Toint, "A globally convergent lagrangian barrier algorithm for optimization with general inequality constraints and simple bounds," *Mathematics of Computation of the American Mathematical Society*, vol. 66, no. 217, pp. 261–288, 1997.
- [34] K. Soomro, A. R. Zamir, and M. Shah, "UCF101: A dataset of 101 human actions classes from videos in the wild," *in Int. Workshop on*

*Action Recognition with a Large Number of Classes in conjunction with Int. Conf. Computer Vision (ICCV), 2013.*

- [35] T. Brox and J. Malik, "Large displacement optical flow: descriptor matching in variational motion estimation," *IEEE Trans. on Patt. Anal. Mach. Intel.*, vol. 33, no. 3, pp. 500–513, Mar 2011.
- [36] Y. Peng, Y. Zhao, and J. Zhang, "Two-stream collaborative learning with spatial-temporal attention for video classification," *IEEE Trans. Circuits and Syst. for Video Technol.*, DOI:10.1109/TCSVT.2018.2808685, to appear.
- [37] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool, "Temporal segment networks: towards good practices for deep action recognition," *Proc. IEEE Euro. Conf. Computer Vision (ECCV)*, pp. 20–36, 2016.
- [38] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3D convolutional networks," *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, pp. 4489–4497, 2015.
- [39] T. Brox, A. Bruhn, N. Papenber, and J. Weickert, "High accuracy optical flow estimation based on a theory for warping," *Proc. IEEE Euro. Conf. Computer Vision (ECCV)*, pp. 25–36, 2004.



**Aaron Chadha** obtained an MA and MEng in Information and Computer Engineering from the University of Cambridge, graduating with a Distinction. He joined the Department of Electronic and Electrical Engineering at University College London in 2014, where he obtained the PhD in the area of image/video classification and domain adaptation under the supervision of Professor Yiannis Andreopoulos. He was awarded with an Industrial Fellowship from the Royal Commission for the Exhibition of 1851 in 2016, which partly funded his work. He is currently a Research Associate with UCL Electronic Engineering.



**Mohammad Jubran** is an Assistant Professor at the Department of Electrical and Computer Engineering at Birzeit University, Birzeit, Palestine. He has earned the B.S. degree in Electronic Engineering from Al-Quds University, Abu-Dies, Palestine, in 2000, and the M.S. and Ph.D. degrees in Electrical Engineering from the University at Buffalo, The State University of New York, Buffalo, NY, USA in 2005 and 2007, respectively. Dr. Jubran was a Visiting Postdoctoral Research Associate at University College London, London, UK (2017-

2018) working on Video Processing, Communication and Classification using Machine Learning. Dr. Jubran is also an active member of IEEE and IEEE Communications Society. He has also served as the counselor of the IEEE student chapter at Birzeit University. Dr. Jubran research interests include video communication systems, wireless and mobile communications and machine learning. He published many articles in this area. He has also participated and led many international projects such as the e-Academy to support Smart Cities Operations in Palestine (eSCO) project funded by EC through the ERAMUS+ program. Dr. Jubran has received many regional and international awards such as the Distinguished Scholar Award - Arab Fund Program, Fulbright Scholarship, Zamalah Fellowship, and Academic Development Award from Birzeit University.



**Yiannis Andreopoulos** obtained the Electrical Engineering Diploma and an MSc from the University of Patras, Greece, and the PhD in Applied Sciences from the Vrije Universiteit Brussel, Belgium. He is a cofounder of iSize Ltd. (London, U.K.), as well as Professor of Data and Signal Processing Systems at the Electronic and Electrical Engineering Department of University College London, London, U.K. His research interests are in machine learning and multimedia systems. In his academic work, he has made major contributions to image/vision computing and large-scale multidimensional data analysis. In collaboration with industry, elements of his research in video processing systems have been integrated into commercial products. Author of 150+ papers, 3 patents, 10+ contributions to standards and 3M+ in research grants as Principal Investigator at UCL.



**Alhabib Abbas** obtained his BSc in Electrical and Electronic Engineering from the University of Bahrain, and his MSc in Telecommunications awarded with Distinction from University College London, U.K. He joined the Communication and Information Systems Group at University College London in 2015 where he is currently a research assistant and is working towards a PhD degree in machine learning and video analysis under the supervision of Professor Yiannis Andreopoulos. His research interests are in machine learning, image and video semantic analysis, and error-tolerant computing in multimedia systems.

video semantic analysis, and error-tolerant computing in multimedia systems.