

Combined Static and Motion Features for Deep-Networks-Based Activity Recognition in Videos

Sameera Ramasinghe¹, *Member, IEEE*, Jathushan Rajasegaran, *Student Member, IEEE*,
Vinoj Jayasundara, *Student Member, IEEE*, Kanchana Ranasinghe, *Student Member, IEEE*,
Ranga Rodrigo, *Member, IEEE*, and Ajith A. Pasqual, *Member, IEEE*

Abstract—Activity recognition in videos in a deep-learning setting—or otherwise—uses both static and pre-computed motion components. The method of combining the two components, while keeping the burden on the deep network less, still remains uninvestigated. Moreover, it is not clear what the level of contribution of individual components is, and how to control the contribution. In this paper, we use a combination of convolutional-neural-network-generated static features and motion features in the form of motion tubes. We propose three schemas for combining static and motion components: based on a variance ratio, principal components, and Cholesky decomposition. The Cholesky-decomposition-based method allows the control of contributions. The ratio given by variance analysis of static and motion features matches well with the experimental optimal ratio used in the Cholesky decomposition-based method. The resulting activity recognition system is better or on par with the existing state-of-the-art when tested with three popular data sets. The findings also enable us to characterize a data set with respect to its richness in motion information.

Index Terms—Activity recognition, fusing features, convolutional neural networks (CNN), recurrent neural networks (RNN), long short-term memory (LSTM).

I. INTRODUCTION

AUTOMATIC activity recognition in videos is an intensely researched area in computer vision due to its wide range of real-world applications in sports, health care, surveillance, robot vision, and human-computer interaction. Furthermore, the rapid growth of digital video data demands automatic classification and indexing of videos. Despite the increased interest, the state-of-the-art systems are still far from human-level performance, in contrast to the success in image classification [1], [2]. This is partially due to the complex intra-class variations in videos, some obvious causes being the view point,

background clutter, high dimensionality of data, lack of large datasets, and low resolution.

Despite these reasons more-or-less affecting automatic image classification, it has been quite successful in recent years, largely owing to the rise of deep learning techniques. This is not the case in video classification, although deep learning is starting to be widely applied. Therefore, it is worthwhile investigating what is holding back video classification. In this study, we address three key areas: exploiting the underlying dynamics of sub-events for high-level action recognition, crafting self-explanatory, independent static and motion features—in which, motion features should capture micro-level actions of each actor or object independently, such as arm or leg movements—and optimum fusing of static and motion features for better accuracy.

A complex activity typically comprises several sub activities. The existing approaches try to classify a video treating it as a single, high-level activity [3]–[6]. As the action becomes complex, the behavior and the temporal evolution of its underlying sub-events become complicated. For example, cooking may involve sub-events: cutting, turning on the cooker, and stirring. It may not always preserve this same lower order for the same action class. Instead, they may contain a higher-order temporal relationship among them. For example, turning on the cooker and cutting may appear in reverse order in another video in the same class. Therefore, this temporal pattern of sub-events is not easy to capture through a simple time series analysis. These patterns can only be identified by observing many examples, using a system which has an infinite dynamic response. Therefore, it is important to model this higher-order temporal progression, and capture temporal dynamics of these sub-events for better recognition of complex activities.

In contrast to image classification, the information contained in videos are not in a single domain. Both the motion patterns of the actors and objects, as well as the static information—such as, background and still objects that the actors interact with—are important for determining an action. For example, the body movements of a group of people fighting may closely relate to the body movements of a sports event, e.g., wrestling. In such a case, it is tough to distinguish between the two activities solely by looking at the motion patterns. Inspecting the

Manuscript received February 1, 2017; revised August 2, 2017; accepted September 21, 2017. Date of publication October 6, 2017; date of current version September 4, 2019. This work was supported by the National Research Council of Sri Lanka under Grant 12-018. This paper was recommended by Associate Editor Y. Li. (*Corresponding author: Sameera Ramasinghe.*)

The authors are with the Department of Electronic and Telecommunication Engineering, University of Moratuwa, Moratuwa 10400, Sri Lanka (e-mail: samramasinghe@gmail.com; brjathu@gmail.com; vinojjayasundara@gmail.com; kahnchana@gmail.com; ranga@uom.lk; pasqual@ent.mrt.ac.lk).

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCSVT.2017.2760858

1051-8215 © 2017 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

background setting and objects is crucial in such a scenario. Therefore, it is necessary to engineer powerful features from both motion and static domains. In addition, these features must be complementary, and one feature domain should not influence the other domain. In other words, they should be self-explanatory, and mutually exclusive, as much as possible. Also, these motion features should be able to capture activities of each actor independently. If so, the distributional properties of these actions, over short and long durations, can be used to identify high level actions. In the proposed method, we craft static and motion features to have this property.

Furthermore, how to optimally combine or fuse these motion and static descriptors remains a challenge for three key reasons: both static and motion information provide cues regarding an action in a video, the contribution ratio from each domain affects the final recognition accuracy, and optimum contribution ratio of static and motion information may depend on the richness of motion information in the video. The combined descriptor should contain the essence of both domains, and should not have a negative influence on each other. Recently, there has been attempts to answer this question [6], [5]. However, these existing methodologies lack the insight in to how much this ratio affects the final accuracy, and has no control over this ratio. One major requirement of the fusion method is to have a high-level intuitive interpretation on how much contribution each domain provides to the final descriptor, and can control the level of contribution. This ability makes it possible to deeply investigate the optimum contribution of each domain for a better accuracy. In this study, we investigate this factor extensively.

In this work, we focus on video activity recognition using both static and motion information. We address three problem areas: crafting mutually exclusive static and motion features, optimal fusion of the crafted features, and modeling temporal dynamics of sub-activities for high-level activity recognition. In order to examine the temporal evolution of sub-activities subsequently, we create video segments with constant overlapping durations. Afterwards, we combine static and motion descriptors to represent each segment. We propose *motion tubes*, a dense trajectory [3] based tracking mechanism, to identify and track candidate moving areas separately. This enables independent modeling of the activities in each moving area. In order to capture motion patterns, we use histogram oriented optic flows (HOOF) [7] inside motion tubes. Then we apply a bag-of-words (BoW) method on the generated features to capture the distributional properties of micro actions, such as body movements, and create high level, discriminative motion features.

Inspired by the power of object recognition of convolutional neural networks (CNNs), we create a seven-layer deep CNN, and pre-train it on the popular ImageNet dataset [8]. Afterwards, we use this trained CNN to create deep features, to synthesize static descriptors. Then, using a computationally efficient, yet powerful, mathematical model, we fuse static and motion feature vectors. We propose three such novel methods in this paper: based on Cholesky decomposition, variance ratio of motion and static vectors, and principal components analysis (PCA). The Cholesky decomposition based model

provides the ability to precisely control the contribution of static and motion domains to the final fused descriptor. Using this intuition, we investigate the optimum contribution of each domain, experimentally. The variance ratio based method also provides us this ability, and additionally, lets us obtain the optimum contribution ratio mathematically. We show that the optimum contribution ratio obtained experimentally using the Cholesky based method matches with the ratio obtained mathematically from the variance based method. Furthermore, we show that this optimum contribution may vary depending on the richness of motion information, and affects the final accuracy significantly.

In order to model the temporal progression of sub events, we feed the fused vectors to a long short-term memory (LSTM) network. The LSTM network discovers the underlying temporal patterns of the sub events, and classifies high level actions. We feed the same vectors to a classifier which does not capture temporal dynamics to show that modeling temporal progression of sub events indeed contributes for a better accuracy. We carry out our experiments on the three popular action data sets, UCF-11 [9], Hollywood2 [10], and HMDB51 [11].

The key contributions of this paper are as follows:

- We propose an end-to-end system, which extracts both static and motion information, fuses, and models the temporal evolution of sub-events, and does action classification.
- We propose a novel, moving actor and object tracking mechanism, called *motion tubes*, which enables the system to track each actor or object independently, and model the motion patterns individually over a long time period. This allows the system to model actions occurring in a video in micro level, and use these learned dynamics at a high level afterwards.
- We propose three novel, simple, and efficient mathematical models for fusing two vectors, in two different domains, based on Cholesky transformation, variance ratio of motion and static vectors, and PCA. The first two methods provide the ability to govern the contribution of each domain for the fused vector precisely and find the optimum contribution of the two domains, mathematically or empirically. Using this advantage, we prove that static and motion information are complementary and vital for activity recognition through experiments.
- We prove that the final recognition accuracy depends on the ratio of contribution of static and motion domains. Also, we show that this optimum ratio depends on the richness of motion information in the video. Hence, it is beneficial to exploit this optimum ratio for a better accuracy.
- We model the underlying temporal evolution of sub-events for complex activity recognition using an LSTM network. We experimentally prove that capturing these dynamics indeed benefits the final accuracy.

With the proposed techniques we outperform the existing best results for the datasets UCF-11 [9] and Hollywood2 [10], and are on par for the dataset HMDB51 [11].

II. RELATED WORK

There has not been many approaches in activity recognition, which highlight the importance of exclusively engineered static and motion features. Most of the work rely on generating spatio-temporal interest regions, such as, action tubes [12], tubelets [13], dense trajectory based methods [14], [15], spatio-temporal extensions of spatial pyramid pooling [16], or spatio-temporal low-level features [3], [17]–[21]. Action tubes [12] is quite similar to our motion tubes, but our motion candidate regions are chosen based on more powerful dense trajectories [3] instead of raw optic flows. Also, we employ a tracking mechanism of each moving area through motion tubes isolating actions of each actor throughout the video. This is an extension of the human-actor tracking presented by Wang [4]. Our static interest regions are independent from motion, unlike in Gkioxari and Malik [12], where can extract background scenery information using CNNs, for action recognition. A common attribute of these methods is that motion density is the dominant factor for identifying candidate regions. In contrast, we treat motion and static features as two independent domains, and eliminate the dominance factor.

A few attempts has recently been made on exclusive crafting and late fusion of motion and static features. Simonyan and Zisserman [5] first decomposes a video in to spatial and temporal components based on RGB and optical flow frames. Then they apply two deep CNNs on these two components separately to extract spatial and temporal information. Each network operates mutually exclusively and performs action classification independently. Afterwards, softmax scores are coalesced by late fusion.

Work done in Feichtenhofer *et al.* [22] is also similar. Instead of late fusion, they fuse the two domains in a convolutional layer. Both these approaches rely explicitly on automatic feature generation in increasingly abstract layers. While this has provided promising results on static feature generation, we argue that motion patterns can be better extracted by hand-crafted features. This is because temporal dynamics extend to a longer motion duration unlike spatial variations. It is not possible to capture and discriminate motion patterns in to classes by a system which has a smaller temporal support. There are models which employ 3D convolution [23], [24], which extends the traditional CNNs into temporal domain. Ramasinghe and Rodrigo [6] apply CNNs on optic flows, and Kim *et al.* [25] on low level hand-crafted inputs (spatio-temporal outer boundaries volumes), to extract motion features. However, even by generating hierarchical features on top of pixel level features, it is not easy to discriminate motion classes as the duration extent is short. Also, tracking and modeling actions of each actor separately in longer time durations is not possible with these approaches. Our motion features, on the other hand, are capable of capturing motion patterns in longer temporal durations. Furthermore, with the aid of *motion tubes* our system tracks and models the activities of each moving area separately.

In the case of work done by Wang *et al.* [26], their use of IDT features as motion descriptors and CNN features as static descriptors serves as a baseline for our work.

Considering their experiments and observations, we focus solely on HOOOF creation as opposed to multiple different descriptors for extracting the motion information. In addition, we look into the extraction of micro actions through our work. Also with regards to the approach for static and motion vector fusion, which involves a constant weighing factor, our work explores alternative approaches with focus on three different methods used across all experimenting. In addition, we improve with regards to capturing the temporal evolution of videos using recurrent neural networks, considering the shortcomings of SVM classifiers in capturing temporal evolutions.

Regarding video evolution, Fernando *et al.* [27] postulate a function capable of ordering the frames of a video temporally. They learn a ranking function per video using a ranking machine and use the learned parameters as a video descriptor. Several other methodologies, e.g., HMM [28], [29], CRF-based methods [30], also have been employed in this aspect. These methods model the video evolution in frame level. In contrast, attempts for temporal ordering of atomic events also has been made [31], [32]. Rohrbach *et al.* [31], encode transition probabilities of a series of events statistically with a HMM model. Bhattacharya *et al.* [32] identify low level actions using dense trajectories and assign concept identity probabilities for each action. They apply a LDS on these generated concept vectors to exploit temporal dynamics of the low level actions. Li *et al.* [33] uses simple dynamical systems [34], [35] to create a dictionary of low-level spatio-temporal attributes. They use these attributes later as a histogram to represent high level actions. Our method too follows a similar approach, as we also generate descriptors for sub-events and then extract temporal progression of these sub-events. However, instead of a simple statistical model, which has a finite dynamic response, we use an LSTM network [36] to capture these dynamics. In action recognition literature, such models are starting to appear. In Ng *et al.* [37] the LSTM network models the dynamics of the CNN activations, and in Donahue *et al.* [38], the LSTM network learns the temporal dynamics of low level features generated by a CNN.

III. METHODOLOGY

A. Overview

This section outlines our approach. The overall methodology is illustrated in Fig. 1.

Our activity classifier classifies video snippets based on their descriptors. In order to compute descriptors, initially, we segment a video into small snippets of 15 frames with a constant frame overlap. Then we carry out feature construction pipelining for each of these snippets, as shown in Fig. 1. We compute features for each snippet that describe both motion and static domains. For extracting motion features, we create *motion tubes* across frames, where we track each moving area across the frames using “action boxes”. Action boxes are square regions, which exhibit significant motion in each frame. We choose candidate areas by first creating dense trajectories for each frame, and then clustering trajectory points preceded by a significant amount of pre-processing. This process is explained in sub-section III-B. These action

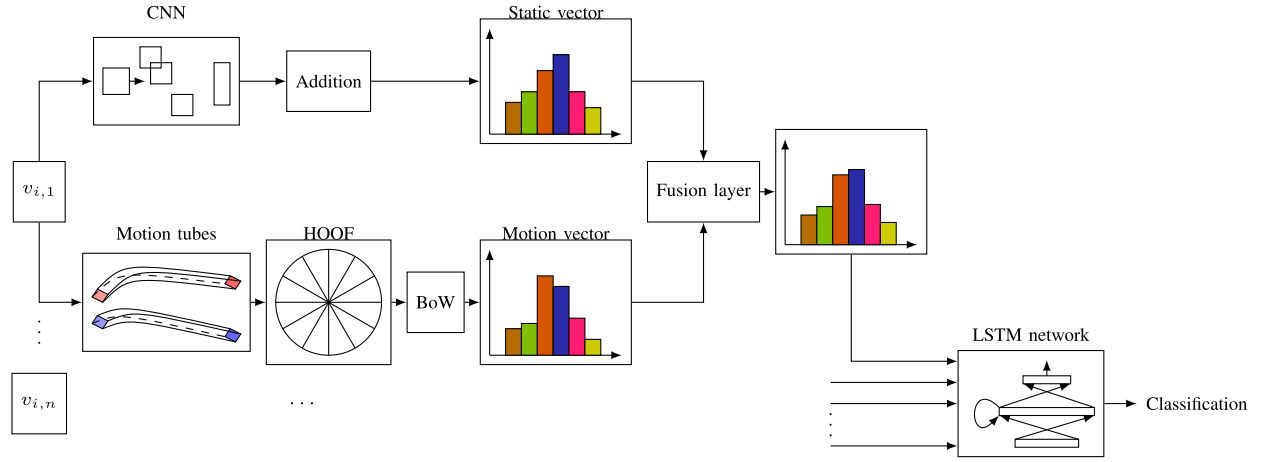


Fig. 1. **Overall methodology.** The whole process consists of five major steps: (i) segmenting a video (ii) crafting static features, (iii) crafting motion features, (iv) fusing static and motion features, and (v) capturing temporal evolution of sub events. Static and motion features are independent and complementary. We generate static features based on a pre-trained CNN and motion features based on motion tubes, and capture the temporal evolution of sub events using an LSTM network.

boxes create motion tubes by getting linking across the frames. Then we calculate HOOF [7] features within these motion tubes and apply a bag-of-features method on these to create a motion descriptor for each video segment.

For extracting static features we train a deep CNN on ImageNet (Fig. 2). Then we apply this CNN on the frames of each video snippet to retrieve deep features—output vector from the final softmax layer of the CNN—from it. Then we use these features to create a static descriptor for the video segment. Afterwards, we combine these motion and static descriptors using one of the fusion models described in subsection III-C.

The system can then represent a video as a vector time series, $C = [c_{t_0}, c_{t_1}, \dots, c_{t_{n-1}}]$, where n is the number of segments. Then we apply an LSTM network on these features and exploit the dynamics of time evolution of the combined vector. Finally, we classify the dynamics of this time series and predict actions.

B. Motion Features

This section discusses the detailed methodology of creating motion features, particularly, “Motion tubes”, “HOOF”, and “BoW” blocks as illustrated in Fig. 1.

1) *Low Level Motion Descriptor*: A pixel-level descriptor is required to identify moving points in the video. Dense trajectories [3] is a powerful pixel-level algorithm, which captures and tracks motion across several frames. In this work, we choose dense trajectories as the low-level descriptor for capturing raw motion.

2) *Clustering*: Initially, we create dense trajectories for every frame in the video. Then in order to isolate each sub area in a frame which contains significant motion we apply DBScan clustering on the calculated trajectory points. Algorithm 1 illustrates our clustering approach. Empirically, we use 8 and 10 as ϵ and MinPoints parameters respectively.

Despite the presence of many regions which contain motion in a video, some are neither significant nor descriptive. Those moving regions can be neglected without loss of performance.

Algorithm 1 DBScan Clustering Algorithm

Require: \mathcal{D} ▷ Dataset of sub-areas in frame
Require: M ▷ Min. no. of points
Require: ϵ ▷ Max. cluster radius
▷ initialize cluster no.

- 1: $c \leftarrow 0$
- 2: **for each** $P \in \mathcal{D}$ **do do**
- 3: **if** P is visited **then**
- 4: **continue**
- 5: **end if**
- 6: mark P as visited
- 7: NeighborPts = GETALLPOINTS(P, ϵ)
- 8: **if** size(NeighborPts) < M **then**
- 9: mark P as noise
- 10: **else**
- 11: $c = \text{next cluster}$
- 12: ADDTOCLUSTER($P, \text{NeighborPts}, c, \epsilon, M$)
- 13: **end if**
- 14: **end for**
- 15: **function** ADDTOCLUSTER($P, \text{NeighborPts}, c, \epsilon, M$)
- 16: add P to cluster c
- 17: **for** doeach point $np \in \text{NeighborPts}$ **do**
- 18: **if** P is not visited **then**
- 19: mark np as visited
- 20: NeighborPts' = GETALLPOINTS(np, ϵ)
- 21: **if** size(NeighborPts) $\geq M$ **then**
- 22: NeighborPts' \leftarrow NeighborPts joined with NeighborPts'
- 23: **end if**
- 24: **end if**
- 25: **if** np is not yet member of any cluster **then**
- 26: add np to cluster c
- 27: **end if**
- 28: **end for**
- 29: **end function**
- 30: **function** GETALLPOINTS(P, ϵ) **return** all points within P 's ϵ -neighborhood (including P)
- 31: **end function**

Therefore, after clustering each trajectory point in to cluster groups non-significant cluster groups are ignored. We do this in order to prevent the algorithm focusing on small random moving areas in a video and to limit creation of motion tubes to areas which are significant and descriptive. Therefore all the clusters which are not at least 50% the size of the largest cluster of the frame are discarded.

After identifying the initial candidate clusters for creating action boxes, further processing is done to each cluster to ensure that it contains only the important moving areas according to Algorithm 2. For each point, the Chebychev distance from the centroid of the cluster is calculated. We discard the furthest 20% of the points from the cluster. The reason behind the choice of Chebychev distance over Euclidean distance is due to the possibility of obtaining symmetric square shaped cluster groups as opposed to circular ones. This makes it easier to track moving areas and create motion tubes.

Algorithm 2 Boundary Noise Removal Algorithm of Clusters

Require: M_d ▷ Max. Chebychev dist.

Require: \mathcal{C} ▷ Input cluster

```

1: totalPoints ← points within  $M_d$  of center of  $\mathcal{C}$ 
2: currentPoints ← totalPoints
3: while true do
4:   if COUNT(currentPoints) < COUNT(totalPoints) × 0.8
     then return currentPoints
5:   end if
6:    $M_d \leftarrow M_d - 1$ 
7:   currentPoints ← points within  $M_d$  of center of  $\mathcal{C}$ 
8: end while

```

After identifying square-shaped interest regions (action boxes), we represent each of them with a vector, $b = (x, y, r, f)$, where x and y are the coordinates of the top left corner of the box, r is the height or width of the box, and f is the frame number.

3) *Motion Tubes*: Since our work models the time evolution of sub activities within a video, we divide each input video V_i into temporal segments, $f(V_i) = [v_{i,1}, v_{i,2}, \dots, v_{i,n}]$, and create features for each individual segment separately. Therefore, after creating the action boxes for each video segment, the action boxes within a segment $v_{i,t}$ can be represented as,

$$g(v_{i,t}) = \{[b_{t,1,1}, b_{t,1,2}, \dots, b_{t,1,q}], \\ [b_{t,2,1}, b_{t,2,2}, \dots, b_{t,2,p}], \\ \dots, [b_{t,n,1}, b_{t,n,2}, \dots, b_{t,n,k}]\} \quad (1)$$

where $b_{t,j,k}$ is the k^{th} action box in j^{th} frame of the t^{th} video segment. Note that the number of action boxes differ from frame to frame.

Therefore, before linking the action boxes, to create motion tubes further pre-processing is needed to ensure the same number of action boxes exist in every frame within a video segment. Otherwise, different motion tubes could become joined halfway through the video, and the effort to capture dynamics of each moving object separately is disturbed.

For this purpose, first we calculate the mean number of action boxes per frame in each segment. Then we obtain the

rounded down value, N , of the mean number. Afterwards, we iterate through each frame starting from frame number 1 until we come to a frame W which has N number of action boxes. Then from frame W we propagate forward and backward along the frames, either to eliminate or add action boxes. The procedure is explained below. If the action box count in a particular frame is larger than the previous frame, the smallest excess number of action boxes are removed.

In the case where the action box count is lower than the previous frame, linear regression is used for each x, y and r value of vector $b = (x, y, r, f)$ up to that frame, in order to create artificial action boxes until the number of action boxes matches N .

Note how this processing results in Eq. 1 being transformed in to Eq. 2, thus verifying that the number of action boxes per frame is equal for all frames within a video segment.

$$h(g(v_{i,t})) = \{[b_{t,1,1}, b_{t,1,2}, \dots, b_{t,1,k}], \\ [b_{t,2,1}, b_{t,2,2}, \dots, b_{t,2,k}], \\ \dots, [b_{t,n,1}, b_{t,n,2}, \dots, b_{t,n,k}]\} \quad (2)$$

The following procedure is used to link the action boxes in consecutive frames. Assume $b_{t,k,1}, b_{t,k,2}, \dots, b_{t,k,n}$ and $b_{t,k+1,1}, b_{t,k+1,2}, \dots, b_{t,k+1,n}$ are action boxes in two consecutive frames at time k and $k+1$. Then the following distance matrix is calculated.

$$D = \begin{bmatrix} D_{11} & D_{12} & D_{13} & \dots & D_{1k} \\ D_{21} & D_{22} & D_{23} & \dots & D_{2k} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ D_{k1} & D_{k2} & D_{k3} & \dots & D_{kk} \end{bmatrix} \quad (3)$$

where $D_{i,j}$ is the Euclidean distance between the centroids of i^{th} action box in k^{th} frame and j^{th} action box in $(k+1)^{st}$ frame. Then u^{th} action box at $k+1$, and 1^{st} action box at k are linked, where u is found using,

$$u = \underset{j \in J}{\operatorname{argmin}} \{D_{1,j}\}, J = \{1, 2, \dots, l\} \quad (4)$$

Then the 1^{st} row and the u^{th} column are removed from the distance matrix, and we apply the same process repeatedly using Eq. 4 to link each of the action boxes at k with $k+1$.

By this removal process, we avoid combining of motion tubes half-way through the video segment and keep them isolated from each other, which is vital for capturing the dynamics separately for each moving object.

Finally, we create a $(z \times n)$ -by-5 matrix M_i — z and n are number-of-frames and number-of-action-boxes-per-frame, respectively—which encodes all the information of motion tubes, in a particular video segment i . The rows of M_i for the k^{th} frame is shown in Eq. 5,

$$M_i = \begin{bmatrix} k & 1 & x_{z,1} & y_{z,1} & r_{z,1} \\ k & 2 & x_{z,2} & y_{z,2} & r_{z,2} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ k & n & x_{z,n} & y_{z,n} & r_{z,n} \end{bmatrix} \quad (5)$$

where the columns represent the frame number, action box number, x coordinate of the top left corner of the action box, y coordinate of the top left corner of the action box, and the width/height of the action box, respectively.

4) *Histogram Oriented Optic Flows (HOOF)*: Since each action box in a particular motion tube may differ in size, we take $R = \max(r_i)$, for $\forall i$, where r_i is the length of the i^{th} action box of the motion tube. Then we redraw the action boxes around their centroids having width or length as R . After identifying the k number of motion tubes (k is a variable) for each video segment $v_{i,n}$, we calculate the optic flows along each motion tube using Lucas *et. al* [39]. After that we create HOOF [7] for every action box within a motion tube. Each optic flow vector within a spatio-temporal action box within a motion tube is binned according to its primary angle from the horizontal axis and weighted according to its magnitude. Therefore, all optical flow vectors, $z = [x, y]^T$ with direction, $\theta = \tan^{-1}(\frac{y}{x})$ in the range,

$$-\frac{\pi}{2} + \pi \frac{b-1}{B} \leq \theta < -\frac{\pi}{2} + \pi \frac{b}{B} \quad (6)$$

will contribute a weight of $\sqrt{x^2 + y^2}$ to the sum in bin b , $1 \leq b \leq B$ out of a total of B bins. Finally, the histogram is normalized. We choose 100 number of bins.

5) *Bag of HOOFs*: We use a bag of features method to create a motion descriptor for each video segment. First, we create a code book for HOOF vectors. 100,000 vectors are randomly selected from all the HOOF vectors of all the video segments in all video classes. Then these 100,000 vectors are clustered using k -means clustering and 1000 cluster heads are identified. We choose the number of cluster heads as 1000, because the dimensions of final motion descriptors are needed to be the same as the static descriptors, which is explained in section V. Then for each video segment $v_{i,n}$, a histogram is calculated as follows.

We calculate,

$$p = \underset{j \in J}{\operatorname{argmin}} (T_j - h_{n,k}), J = \{1, 2, \dots, 1000\} \quad (7)$$

for each k in $\{1, 2, \dots, l\}$, where $h_{n,k}$ is the k^{th} HOOF vector of the n^{th} video segment, and T_j is the j^{th} cluster head. Then we increment the histogram values as,

$$H_n(p) = H_n(p) + 1 \quad (8)$$

where $H_n(p)$ is the p^{th} value, $1 \leq p \leq 1000$, of histogram of the n^{th} video segment $v_{i,n}$. After calculating the histogram vector H_n for every video segment $v_{i,n}$ this $H = [H_1, H_2, \dots, H_n]$ is the vector time series, which encodes the time evolution of motion information in the video.

C. Fusing of Static and Motion Features

This work depends on three factors; both static and motion information are vital for action recognition, but the final accuracy depends on the ratio of contribution of each domain, and the optimum contribution may depend on the richness of motion information in the video. We derive our fusion models addressing all these aspects. The three mathematical models we use to fuse the static and motion vectors are described next. This sub section relates to the ‘‘Fusion layer’’ block in Fig. 1.

1) *Cholesky Transformation Based Method*: This derivation is based on the Cholesky transformation. An abstract version of Cholesky transformation is described below.

Let P and Q be two random variables of unknown correlation. These random variables can be transformed into two new random variables (R and S) with a known correlation of ρ , where the value of ρ can be chosen at will. The transformation can be performed as follows.

$$\begin{bmatrix} Y \\ Z \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \rho & \sqrt{1-\rho^2} \end{bmatrix} \times \begin{bmatrix} P \\ Q \end{bmatrix} \quad (9)$$

Therefore,

$$Y = P \quad (10)$$

and

$$Z = \rho P + \sqrt{1-\rho^2} Q \quad (11)$$

The Cholesky transformation guarantees that the correlation between the two random variables Y and Z is ρ .

Based on the above properties of the Cholesky transformation, we propose the following methodology to fuse the static and motion vectors.

Let S and M be static and motion vectors, respectively. Cholesky transformation can be applied to the two vectors S and M with the correlation value ρ_1 .

$$\begin{bmatrix} Y \\ Z \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \rho_1 & \sqrt{1-\rho_1^2} \end{bmatrix} \times \begin{bmatrix} S \\ M \end{bmatrix} \quad (12)$$

$$Y = S \quad (13)$$

$$Z = \rho_1 S + \sqrt{1-\rho_1^2} M \quad (14)$$

Similarly, the transformation can be applied to M and S with the correlation value ρ_2 .

$$\begin{bmatrix} A \\ B \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \rho_2 & \sqrt{1-\rho_2^2} \end{bmatrix} \times \begin{bmatrix} M \\ S \end{bmatrix} \quad (15)$$

$$A = M \quad (16)$$

$$B = \rho_2 M + \sqrt{1-\rho_2^2} S \quad (17)$$

Again, the Cholesky transformation guarantees the following two properties.

- 1) The correlation between S and Z is ρ_1 .
- 2) The correlation between M and B is ρ_2 .

Therefore, if the values of ρ_1 and ρ_2 are chosen in such a way that they obey the following rule,

$$\rho_2 = \sqrt{1-\rho_1^2} \quad (18)$$

it can be guaranteed that $Z = B$, $\forall S, M, \rho_1, \rho_2$. Hence, the resultant vector C can be obtained by,

$$C = Z = B \quad (19)$$

where the correlation between C and S is ρ_1 and the correlation between C and M is ρ_1 . Here S and M represent the static and the motion vectors whereas C represents the resultant vector. This derivation leads us to an important intuition: by choosing the value of ρ_1 , we can choose the degree to which the static features and the motion features contribute,

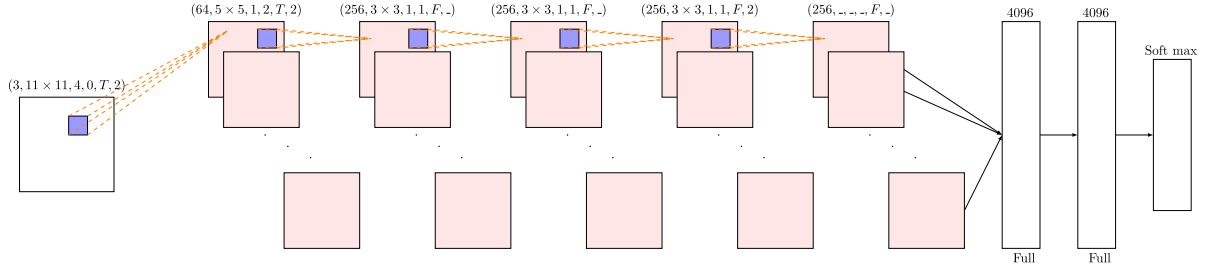


Fig. 2. CNN architecture used for generating static features. The CNN consists of five convolutional layers, two fully connected layers, and one softmax layer. The details of each convolutional layer are provided on top of each layer according to the following format:(number of convolution layers \times filter width \times filter height, convolution stride, spatial padding, is Local Response Normalization added, max-pooling factor). Value above fully connected layers indicates the dimensionality of the layer. We use ReLu as the activation function.

in deriving the resultant vector. In section 4, it is shown, how this property is used to explore, the optimal contribution of static and motion domain information for recognizing actions. The derivation of an optimum ratio between these would require either a continuous variation of the ratio for test datasets, or a theoretical derivation of an optimum ratio. However, we hypothesize that this ratio would depend on various characteristics of a given dataset (which we explore over our experiments), and the intention of this study is towards establishing the existence of an optimum ratio, and not the derivation of that optimum ratio.

2) *Variance Ratio Based Method*: The second method employed to combine the motion and static vectors is based on a Gaussian probability model. We model each vector as a histogram. Using the histogram model, the mean and variance of each vector are calculated in order to fit the two vectors into Gaussian distributions. The joint Gaussian distribution is then computed based on this data.

$$G_{sm}(N) = \frac{1}{2\pi \sigma_m \sigma_s} e^{-\left[\frac{[N-\mu_s]^2}{2\sigma_s^2} + \frac{[N-\mu_m]^2}{2\sigma_m^2}\right]} \quad (20)$$

This computation corresponds to the evaluation line obtained when equating the two random variables. This evaluation line is the diagonal through the origin of the static vector vs motion vector plot. The combined distribution is then obtained through this process.

It must be noted that both histograms (corresponding to static and motion components) do not contain equal information. Hence varying the contribution of each histogram to the resultant distribution is necessary. This requires varying of the evaluation line which can be achieved through scaling of the motion and static axes. This scaling process is carried out by the following matrix.

$$\text{Scaling matrix} = \begin{bmatrix} \frac{\sigma_s}{\sigma_s + \sigma_m} & 0 \\ 0 & \frac{\sigma_m}{\sigma_s + \sigma_m} \end{bmatrix}$$

We may conclude that higher variance of a component along one axis reflects lower detail in the model with regards to the other axis. Considering the motion axis, the contribution of this vector towards the resultant vector may be defined by $(1 - \frac{\sigma_m}{\sigma_m + \sigma_s})$. A high variance always corresponds to a flatter histogram containing less detail.

This parameter we derive is significant as it defines the contribution of each individual motion and static vector pair independent of explicit terms. Hence the optimum ratio for combination of motion and static components of a given dataset can be mathematically evaluated. With regards to the datasets used for experimenting, 30% of motion vector and 70% of static vector constitute this parameter on average. This mathematical inference is further verified through the experiment results in section IV.

Defining new parameters N_s and N_m as follows, we build a new distribution which is a scaled version of the joint Gaussian distribution obtained previously.

$$N_s = N \frac{\sigma_m}{\sigma_m + \sigma_s}$$

$$N_m = N \frac{\sigma_s}{\sigma_m + \sigma_s}$$

Finally, we define the following distribution representative of the combined vector.

$$G_{sm}(N) = \frac{e^{-\frac{1}{2(1-\rho^2)} \left[\frac{[N_s - \mu'_s]^2}{2\sigma_s'^2} + \frac{[N_m - \mu'_m]^2}{2\sigma_m'^2} - \frac{2\rho[N_s - \mu'_s][N_m - \mu'_m]}{\sigma_s'\sigma_m'} \right]}}{2\pi\sigma_m'\sigma_s'\sqrt{1-\rho^2}} \quad (21)$$

The corresponding mean and variance of this newly computed distribution are denoted by μ'_s , σ_s' and μ'_m , σ_m' for the static and motion vector respectively.

3) *PCA Based Method*: The third fusion method is based on Principle Component Analysis (PCA). If there are n number of features in the input vector of the PCA, the output of the PCA will provide a new set of n features which are orthogonal and uncorrelated. Also, if output = $[a_1, a_2, \dots, a_n]$, then $\text{var}(a_1) > \text{var}(a_2) > \dots > \text{var}(a_n)$. Due to the properties of the PCA, the original set of n features can be represented precisely using the first k ($n > k$) principal components of the output vector, given that the total squared reconstruction error is minimized.

In other words, the essence of the original n dimensional dataset is now almost completely represented by the new k dimensional dataset with a minor data loss. Thus, the dimension of the dataset is reduced from n to k .

In our work, the dimensionality of the dataset T is 2 (number of feature domains: static and motion). Using PCA on T , we receive a new set T' , with 2 new features domains.

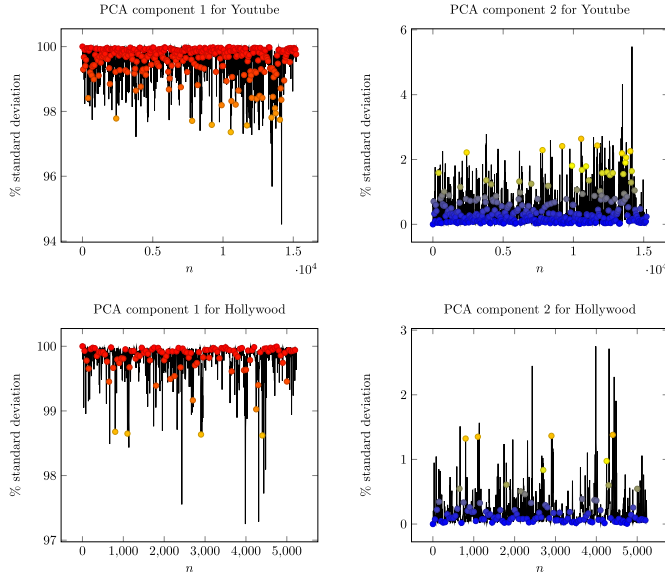


Fig. 3. Percent standard deviation values for the first and second components of the PCA. n is the feature vector index.

We need only one new feature domain in the feature space to represent motion and static domains. Therefore, it is our aim to extract only the first principal component. In order to do this, we need to justify that the first principal component contains a significant majority of the essence of the original dataset. In other words, only a negligible amount of data is lost by eliminating the second principal component.

Therefore, we perform PCA on over 15,000 samples of motion and static vectors and plot the variance percentage of the total variance explained by the first and the second principal components respectively, as in figure and figure

It is evident from the figure that the first principal component almost always accounts for over 97% of the essence of the original dataset, except for only a negligible amount of samples. The lowest percentage registered is approximately 85%, which is still a significantly high value. Figure 3 shows the percent standard deviation values for the first and second components of the PCA for the two datasets.

Therefore, by eliminating the second principal component, only less than 5% of data is lost in average, and hence only the first principal component can be used to accurately represent the original dataset.

D. Capturing Temporal Evolution

Our work requires analyzing complex dynamic temporal patterns in the generated sequences. This sub section relates to the “LSTM network” block shown in Fig. 1 used for this purpose.

In our work, we represent each video as n fixed-length segments (n differs for videos of different lengths) with overlapping frames. Each segment is represented with a fused vector c_t . Therefore, each video can be represented as a vector time series.

Now each vector time series could be analyzed using traditional time series modeling techniques, such as Auto

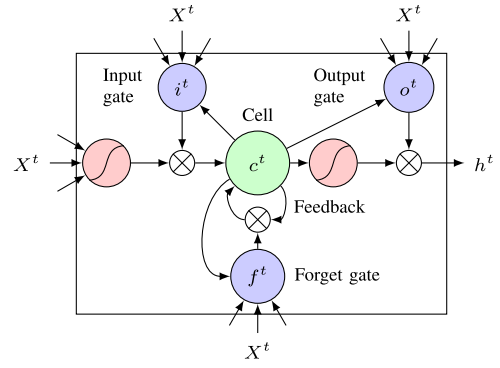


Fig. 4. Long short-term memory (LSTM) block cell. Source [40].

Regressive Moving Average, to obtain features or model parameters that can describe the vector time series. But the main drawback of these methods is that they model current values of a series as a function of past values and have finite dynamic response to time series input. Also, they lack the ability to grasp the internal state representations of a complex time series. RNNs maintain hidden layers with directed feedback connections, and hence, have an infinite dynamic response. While training, it learns internal states of a sequence and usually performs better in modeling complex dynamic temporal patterns of long sequences.

However, it is not ideal to train standard RNNs to solve problems, which require learning of long-term temporal dependencies. This is because of the vanishing-gradient problem, which occurs due to the exponential decay of gradient loss of the function with respect to time.

In practice, LSTM networks typically perform better in such cases. LSTM networks are a special type of RNN, which include a “memory cell”, and as the name suggests, it can maintain a certain state in memory for longer periods of time. It also has a set of control gates for controlling the removal or addition of information to the cell state. This special architecture gives them the ability to capture more long-term dependencies. First, we revise the operation of an LSTM network.

The most important structure of an LSTM unit is its memory cell c_t , which preserves the state. Basic structure of an LSTM unit is shown in figure 4. The memory cell is self-connected, and it has three gates (multiplicative units), i.e., input gate, forget gate and output gate, which are used to control how much long range contextual information of a temporal sequence to store, remove or output.

The detailed activation process of the memory cell and three gates, as shown in Fig. 4 is illustrated as follows:

$$i^t = \sigma(W_{xi}x^t + W_{hi}h^{t-1} + W_{ci}c^{t-1} + b_i) \quad (22)$$

$$f^t = \sigma(W_{xf}x^t + W_{hf}h^{t-1} + W_{cf}c^{t-1} + b_f) \quad (23)$$

$$c^t = f^t c^{t-1} + i^t \tanh(W_{xc}x^t + W_{hc}h^{t-1} + b_c) \quad (24)$$

$$o^t = \sigma(W_{xo}x^t + W_{ho}h^{t-1} + W_{co}c^{t-1} + b_o) \quad (25)$$

$$h^t = o^t \tanh(c^t) \quad (26)$$

where W is the connection weight between two units and $\sigma(\cdot)$ is the sigmoid function.

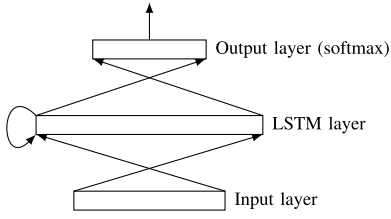


Fig. 5. A simple illustration of the LSTM network. The network consists of an input layer, a 128-unit LSTM layer with 0.8 dropout, and a fully-connected softmax output layer.

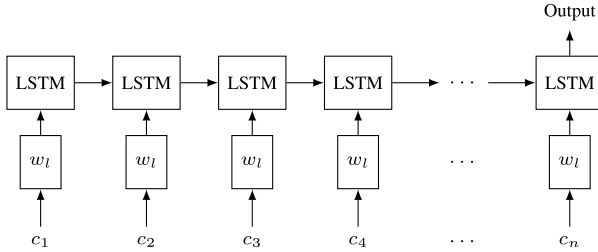


Fig. 6. The process of feeding fused vectors to the LSTM network. c_i indicates the fused vector representing the i_{th} video segment.

Since the LSTM network is used only for capturing the temporal dynamic patterns between sub actions, one LSTM layer is enough. Our LSTM network is shown in Fig. 5. The network consists of an input layer, a 128-unit LSTM layer with 0.8 dropout, and a fully connected softmax output layer. As we have a sequence of activities per classification, we use a many-to-one approach for feeding the fused vectors to the network, as shown in Fig. 6.

IV. EXPERIMENTS AND RESULTS

This section details our experimental methodology and the video datasets used. We evaluate our approach on the two popular datasets UCF-11 and Hollywood2. On both datasets, we show that our work exceeds the current state-of-the-art results (Table II). We also vary the contribution of static and motion features for the calculation of combined vector series and explore what is optimum contribution from each domain. We show that optimum contribution may vary depending on the dataset. We also show that static and motion features are complementary, and provide vital information about the actions occurring in a video. We compare our three fusion models and show that all the methods are better or on par with existing state-of-the-art. Furthermore, we highlight the importance of considering the time evolution of sub activities in order to identify complex events by comparing the results of LSTM and Random Forest Classification algorithm(which does not capture the temporal dynamics), when applied on our features.

A. Datasets

1) *Hollywood 2* [10]: This consists of 12 classes of human actions distributed over 1500 video clips: *answer phone, drive car, eat, fight person, get out car, hand shake, hug person, kiss, run, sit down, sit up, and stand up*. The dataset is composed of

TABLE I
DERIVATION OF ρ VALUES FOR DIFFERENT CONTRIBUTION LEVELS OF STATIC AND MOTION DOMAINS TO THE FUSED VECTOR

| Contribution to Z | ρ value | Fusion vector |
|---|---|---|
| 80% Motion, 20% Static $\rho_1 = 4\rho_2$ | $\frac{1}{4}\rho_1 = \sqrt{1 - \rho_1^2}$ $\rho_1 = \frac{4}{\sqrt{17}}$ | $Z = \frac{4}{\sqrt{17}}M + \frac{1}{\sqrt{17}}S$ |
| 60% Motion, 40% Static $2\rho_1 = 3\rho_2$ | $\frac{2}{3}\rho_1 = \sqrt{1 - \rho_1^2}$ $\rho_1 = \frac{3}{\sqrt{13}}$ | $Z = \frac{3}{\sqrt{13}}M + \frac{2}{\sqrt{13}}S$ |
| 50% Motion, 50% Static $\rho_1 = \rho_2$ | $\rho_1 = \sqrt{1 - \rho_1^2}$ $\rho_1 = \frac{1}{\sqrt{2}}$ | $Z = \frac{1}{\sqrt{2}}M + \frac{1}{\sqrt{2}}S$ |
| 40% Motion, 60% Static $3\rho_1 = 2\rho_2$ | $\frac{3}{2}\rho_1 = \sqrt{1 - \rho_1^2}$ $\rho_1 = \frac{2}{\sqrt{13}}$ | $Z = \frac{2}{\sqrt{13}}M + \frac{3}{\sqrt{13}}S$ |
| 80% Motion, 20% Static $4\rho_1 = \rho_2$ | $4\rho_1 = \sqrt{1 - \rho_1^2}$ $\rho_1 = \frac{1}{\sqrt{17}}$ | $Z = \frac{1}{\sqrt{17}}M + \frac{4}{\sqrt{17}}S$ |

TABLE II
OVERALL ACCURACY OF UCF-11, HOLLYWOOD2, AND HMDB51 FOR VARYING RATIOS BETWEEN STATIC AND MOTION COMPONENTS. THE VECTORS ARE FUSED USING CHOLESKY METHOD. RATIOS ARE INDICATED IN THE FORMAT STATIC:MOTION

| Ratio | UCF-11[9] | Hollywood2[10] | HMDB51-[11] |
|-------|--------------|----------------|---------------|
| 100:0 | 91.8% | 56.9% | 48.2% |
| 80:20 | 96.3% | 80.9% | 62.25% |
| 60:40 | 95.3% | 73.6% | 67.24% |
| 50:50 | 95.3% | 64.9% | 58.64% |
| 40:60 | 93.6% | 60.3% | 42.48% |
| 20:80 | 91.8% | 51.9% | 40.43% |

video clips from 69 movies and provides a challenging task, in automatic action detection.

2) *UCF-11* [9]: This consists over 1000 sports and home videos from YouTube. This dataset contains 11 action classes: *basketball shooting, cycle, dive, golf swing, horse back ride, soccer juggle, swing, tennis swing, trampoline jump, volleyball spike, and walk with a dog*. Each of the action sets is subdivided into 25 groups sharing similar environment conditions. This is a challenging dataset with camera jitter, cluttered backgrounds and variable illumination.

3) *HMDB-51* [11]: This dataset consists of 6849 videos divided into 51 action classes, with each class containing a minimum of 101 videos. The actions categories can be grouped in five types: *general facial actions like smile, laugh, chew, and talk; facial actions with object manipulation like smoke, eat, and drink; general body movements like cartwheel, clap hands, climb, climb stairs; body movements with object interaction like brush hair, catch, draw sword, dribble; and body movements for human interaction like fencing, hug, kick someone, kiss*. This is also a challenging dataset as the use of video clips extracted from real-world videos possess the presence of significant camera/background motion alongside varying illumination.

B. Contribution of Static and Motion Domains

The derivation done in Cholesky based method for fusing the static and motion vectors, provides us an insightful intuition: we can control the contribution of motion and static

TABLE III

PER-CLASS ACCURACY FOR DIFFERENT CONTRIBUTION OF STATIC AND MOTION VECTORS FOR UCF-11. THE VECTORS ARE FUSED USING CHOLESKY METHOD. RATIOS ARE INDICATED IN THE FORMAT STATIC:MOTION. HIGHEST ACCURACY FOR UCF-11 IS ACHIEVED USING A 80:20 RATIO BETWEEN STATIC AND MOTION VECTORS

| Class | 100:0 | 80:20 | 60:40 | 50:50 | 40:60 | 20:80 | 0:100 |
|------------|-------|--------------|-------|--------------|-------|-------|--------|
| B_shooting | 92.4% | 96.3% | 92.7% | 96.3% | 91.3% | 91.9% | 91.3% |
| Biking | 94.3% | 97.8% | 95.6% | 95.4% | 95.4% | 92.6% | 89.5% |
| Diving | 90.3% | 95.8% | 94.3% | 94.3% | 93.1% | 89.6% | 86.2% |
| G_swinging | 93.2% | 96.7% | 96.0% | 95.8% | 93.3% | 92.8% | 90.5% |
| H_riding | 94.0% | 98.0% | 96.6% | 95.6% | 93.1% | 90.2% | 87.2% |
| S_juggling | 92.4% | 96.5% | 96.0% | 96.0% | 93.7% | 90.2% | 85.4% |
| Swinging | 89.3% | 94.3% | 94.3% | 93.6% | 94.1% | 91.7% | 88.2% |
| T_swinging | 92.3% | 96.9% | 95.7% | 94.5% | 94.1% | 93.3% | 90.6% |
| T_jumping | 93.7% | 97.6% | 96.7% | 94.5% | 94.1% | 93.1% | 90.6% |
| V_spiking | 88.2% | 93.4% | 94.2% | 97.2% | 94.1% | 93.0% | 89.3% |
| W_dog | 90.2% | 96.7% | 96.2% | 95.4% | 93.3% | 91.9% | 87.2% |
| Accuracy | 91.8% | 96.3% | 95.3% | 95.3% | 93.6% | 91.8% | 88.72% |

TABLE IV

MAP FOR EACH CLASS FOR DIFFERENT CONTRIBUTION OF STATIC AND MOTION VECTORS TO THE FUSED VECTOR FOR HOLLYWOOD2. RATIOS ARE INDICATED IN THE FORMAT STATIC:MOTION. HIGHEST MAP FOR HOLLYWOOD2 IS ACHIEVED USING A 80:20 RATIO BETWEEN STATIC AND MOTION VECTORS

| Class | 100:0 | 80:20 | 60:40 | 50:50 | 40:60 | 20:80 | 0:100 |
|-------------|-------|--------------|--------------|-------|-------|-------|-------|
| AnswerPhone | 52.3% | 76.6% | 49.6% | 42.4% | 38.2% | 36.6% | 28.2% |
| DriveCar | 54.6% | 98.1% | 49.2% | 42.5% | 39.1% | 37.7% | 26.4% |
| Eat | 50.0% | 62.1% | 55.6% | 53.2% | 53.2% | 50.1% | 40.0% |
| FightPerson | 72.2% | 94.3% | 80.2% | 72.8% | 66.6% | 57.4% | 42.2% |
| GetOutCar | 56.9% | 77.4% | 56.2% | 50.2% | 47.3% | 40.2% | 35.5% |
| HandShake | 42.2% | 78.9% | 80.2% | 72.7% | 64.4% | 50.3% | 42.7% |
| HugPerson | 49.9% | 77.1% | 64.3% | 62.6% | 57.2% | 50.9% | 40.6% |
| Kiss | 49.9% | 85.3% | 86.4% | 70.2% | 68.7% | 60.8% | 45.5% |
| Run | 60.2% | 78.2% | 94.8% | 88.2% | 82.3% | 72.2% | 64.5% |
| SitDown | 80.2% | 86.2% | 91.6% | 80.4% | 76.9% | 67.3% | 56.9% |
| SitUp | 58.7% | 75.0% | 78.2% | 70.2% | 67.4% | 51.8% | 47.2% |
| StandUp | 55.5% | 81.2% | 97.4% | 73.6% | 62.1% | 48.3% | 32.3% |
| mAP | 56.9% | 80.9% | 73.6% | 64.9% | 60.3% | 51.9% | 41.8% |

domains to the fusion vector by varying the ρ value. The derivation of ρ values for different contribution ratios is illustrated in table I.

Results for these different contribution values for UCF-11 and Hollywood2 datasets, are shown in table III and table IV. We use accuracy and mean average precision as performance metrics, for UCF-11 and Hollywood2, respectively. For both datasets, we obtain the optimum contribution ratio as 80:20 between static and motion vectors. In the case of HMDB51 dataset, the optimum contribution ratio is obtained as 60:40 between static and motion vectors.

An overview distribution of the overall performance over different contribution levels, from static and motion domains, for both datasets is shown in Fig. 7 and Fig. 8. We can see that the performance change for different contribution percentages of motion and static domain. Also, the optimum contribution may change depending on the nature of the action and richness of motion or static information in the video. For example, if the motion patterns are indistinguishable across actions, static information plays a critical role, for determining the action, and vice versa. The amount of interaction with objects also may play a key role in determining this ratio. In Fig. 8 it is evident that for action classes which does not highly interact with external objects—such as *kiss*, *run*, *sitdown*, *situp*, *standup* *handshake*—, the optimum

motion:static ratio is 40:60. For other action classes, which interact with objects, optimum motion:static ratio is 20:80. This highlights our hypothesis, that being able to control this contribution explicitly, is vital for an action recognition system.

C. Mathematical Validation of Optimum Contribution

As it is evident from the results of table III and table IV, we experimentally obtain the optimum contribution ratio for the datasets UCF-11 and Hollywood2 as 80:20. In section III, in the derivation of the variance ratio based fusion model, we mathematically obtained values for the optimum contribution as 70:30 for the same datasets. It should be noted that these values closely represent the experimental values, and hence, the results are further verified.

D. Comparison of Fusion Models

The per-class accuracies obtained for each fusion model is illustrated in table V and table VI. Although all three methods give impressive results, Cholesky based fusion model is superior, and has an overall accuracy of 96.3% for UCF-11. For Hollywood2 dataset, it achieves a mean average precision of 80.9%.

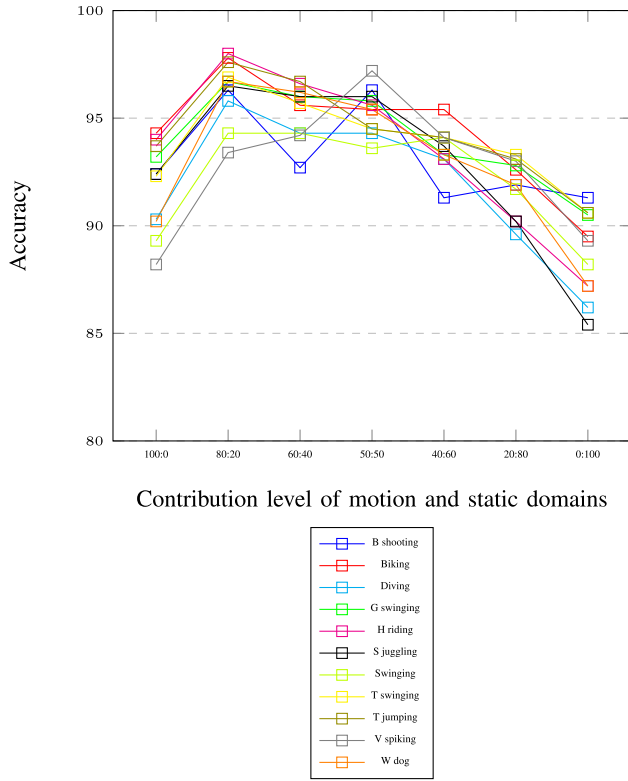


Fig. 7. Accuracy distribution for different contribution levels of motion and static domains. This figure illustrates that motion:static ratio affects the accuracy and the optimum contribution depends on the UCF-11 dataset for each class.

TABLE V

COMPARISON OF FUSION MODELS ON UCF-11 DATASET

| Class | Cholesky | Variance ratio | PCA |
|------------|--------------|----------------|-------|
| B_shooting | 96.3% | 90.3% | 90.6% |
| Biking | 97.8% | 90.8% | 91.0% |
| Diving | 95.8% | 92.3% | 89.3% |
| G_swinging | 96.7% | 90.3% | 92.3% |
| H_riding | 98.0% | 87.4% | 88.6% |
| S_juggling | 96.5% | 89.7% | 92.8% |
| Swinging | 94.3% | 90.0% | 88.0% |
| T_swinging | 96.9% | 89.4% | 93.0% |
| T_jumping | 97.6% | 92.5% | 91.0% |
| V_spiking | 93.4% | 91.6% | 91.7% |
| W_dog | 96.7% | 91.6% | 93.4% |
| Accuracy | 96.3% | 90.5% | 91.1% |

E. Comparison With the State-of-the-Art

Table VII compares our results to state of the art. We use a motion:static ratio of 20:80 for both “change/HMDB” datasets to combine the static and motion vectors, since these values gave the best results. On UCF-11, we significantly outperform the state of the art Ramasinghe and Rodrigo [6] by 3.2%. A mean average precision of 80.9% is achieved by our system for Hollywood2, which outperforms the state-of-the-art by 16.6%. “add here”

Per-action class results, are also compared in table VIII and table IX. In UCF-11, our method excels in 8 out of 11 classes, when compared with Wang *et al.* [3], Ramasinghe and Rodrigo [6], Lucas *et al.* [39], and Ikizler *et al.* [42].

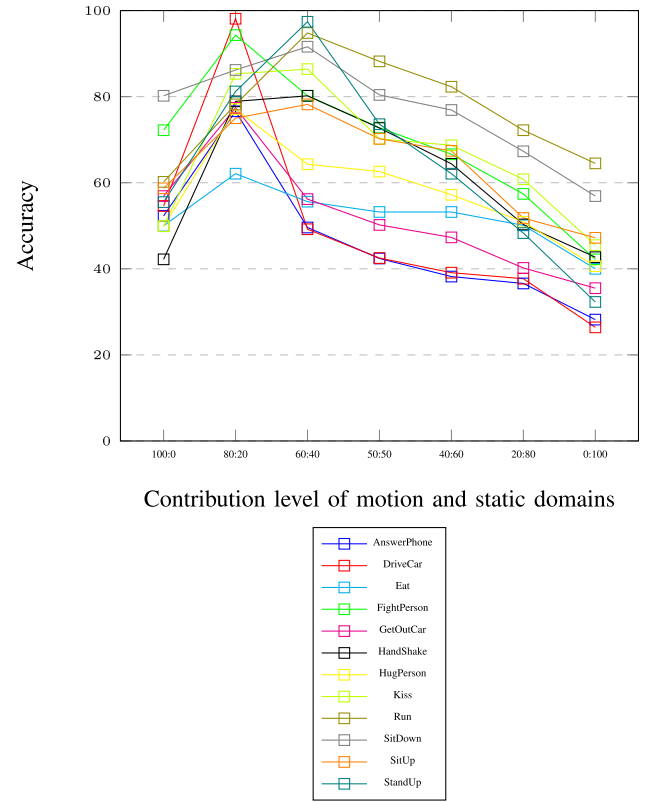


Fig. 8. Accuracy distribution for different contribution levels of motion and static domains. This figure illustrates that motion:static ratio affects the accuracy and the optimum contribution depends on the Hollywood2 dataset.

TABLE VI

COMPARISON OF FUSION MODELS ON HOLLYWOOD2 DATASET

| Class | Cholesky | Variance ratio | PCA |
|-------------|--------------|----------------|-------|
| AnswerPhone | 76.6% | 62.4% | 67.6% |
| DriveCar | 98.1% | 72.8% | 70.0% |
| Eat | 62.1% | 49.4% | 56.5% |
| FightPerson | 94.3% | 78.2% | 72.6% |
| GetOutCar | 77.4% | 46.9% | 56.7% |
| HandShake | 78.9% | 56.9% | 55.6% |
| HugPerson | 77.1% | 52.4% | 60.6% |
| Kiss | 85.3% | 64.0% | 66.6% |
| Run | 78.2% | 58.3% | 54.3% |
| SitDown | 86.2% | 72.0% | 68.6% |
| SitUp | 75.0% | 50.0% | 54.7% |
| StandUp | 81.2% | 54.4% | 50.0% |
| mAP | 80.9% | 59.8% | 61.1% |

In Hollywood2, we calculate the average precision of each class, and compare with Wang *et al.* [3], Lucas *et al.* [39], and Ullah *et al.* [47]. We achieve best results in 10 out of 12 classes in this case.

F. Effectiveness of Capturing Time Evolution

As discussed in earlier sections, complex actions are composed of sub activities preserving a temporal pattern. In this work, we try to capture those underlying patterns by an LSTM network. It is interesting to verify whether this strategy has an impact on the accuracy of the classification. Here we directly feed the fused vectors to a random forest classifier, which does

TABLE VII

COMPARISON OF OUR METHOD WITH STATE-OF-THE-ART METHODS IN THE LITERATURE. STATIC:MOTION RATIOS ARE 80:20 FOR UCF-11 AND HOLLYWOOD2, AND 60:40 FOR HMDB51. THE RESULTS MENTIONED OF OUR SYSTEM ARE THOSE OBTAINED USING THE CHOLESKY METHOD

| UCF-11 | | Hollywood2 | | HMDB51 | |
|-----------------------------------|--------------|--------------------------|--------------|-------------------------|--------------|
| Liu <i>et al.</i> [9] | 71.2% | Vig <i>et al.</i> [41] | 59.4% | Wang <i>et al.</i> [4] | 57.2% |
| Ikizler-Cinbis <i>et al.</i> [42] | 75.21% | Jiang <i>et al.</i> [43] | 59.5% | Wang <i>et al.</i> [15] | 65.9% |
| Wang <i>et al.</i> [3] | 84.2% | Mathe <i>et al.</i> [44] | 61.0% | Zhu <i>et al.</i> [45] | 68.2% |
| Sameera <i>et al.</i> [6] | 93.1% | Jain <i>et al.</i> [46] | 62.5% | | |
| | | Wang <i>et al.</i> [3] | 58.3% | | |
| | | Wang <i>et al.</i> [4] | 64.3% | | |
| Our method | 96.3% | Our method | 80.9% | Our method | 67.24% |

TABLE VIII

PER-CLASS ACCURACY COMPARISON WITH STATE-OF-THE-ART ON UCF-11

| Class | Ours(Cholesky) | KLT[39] | Wang <i>et al.</i> [3] | Ikizler-Cinbis[42] | Ramasinghe <i>et. al.</i> [6] |
|------------|----------------|--------------|------------------------|--------------------|-------------------------------|
| B_shooting | 96.3% | 34.0% | 43.0% | 48.5% | 95.6% |
| Biking | 97.8% | 87.6% | 91.7% | 75.17% | 93.1% |
| Diving | 95.8% | 99.0% | 99.0% | 95.0% | 92.8% |
| G_swinging | 96.7% | 95.0% | 97.0% | 95.0% | 95.0% |
| H_riding | 98.0% | 76.0% | 85.0% | 73.0% | 94.3% |
| S_juggling | 96.5% | 65.0% | 76.0% | 53.0% | 87.8% |
| S_winging | 94.3% | 86.0% | 88.0% | 66.0% | 92.4% |
| T_swinging | 96.9% | 71.0% | 71.0% | 77.0% | 94.9% |
| T_jumping | 97.6% | 93.0% | 94.0% | 93.0% | 94.0% |
| V_spiking | 93.4% | 96.0% | 95.0% | 85.0% | 93.2% |
| W_dog | 96.7% | 76.4% | 87.0% | 66.7% | 91.4% |
| Accuracy | 96.3% | 79.0% | 84.2% | 75.2% | 93.1% |

TABLE IX

PER-CLASS MAP COMPARISON WITH STATE-OF-THE-ART ON HOLLYWOOD2

| Class | Ours | KLT[39] | Wang <i>et al.</i> [3] | Ullah[47] |
|-------------|--------------|---------|------------------------|-----------|
| AnswerPhone | 76.6% | 18.3% | 32.6% | 25.9% |
| DriveCar | 98.1% | 88.8% | 88.0% | 85.9% |
| Eat | 62.1% | 73.4% | 65.2% | 56.4% |
| FightPerson | 94.3% | 74.2% | 81.4% | 74.9% |
| GetOutCar | 77.4% | 47.9% | 52.7% | 44.0% |
| HandShake | 78.9% | 18.4% | 29.6% | 29.7% |
| HugPerson | 77.1% | 42.6% | 54.2% | 46.1% |
| Kiss | 85.3% | 65.0% | 65.8% | 55.0% |
| Run | 78.2% | 76.3% | 82.1% | 69.4% |
| SitDown | 86.2% | 59.0% | 62.5% | 58.9% |
| SitUp | 75.0% | 27.7% | 20.0% | 18.4% |
| StandUp | 81.2% | 63.4% | 65.2% | 57.4% |
| mAP | 80.9% | 54.6% | 58.3% | 51.8% |

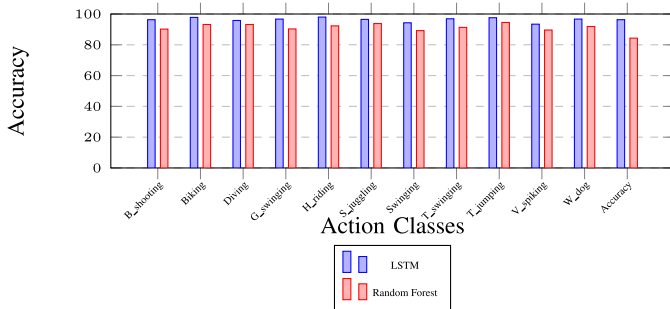


Fig. 9. Accuracy comparison between Random Forest Classifier and LSTM for UCF-11 dataset. Motion:static ratio of 20:80 is used. Accuracy is significantly higher when the temporal dynamics of sub events are captured.

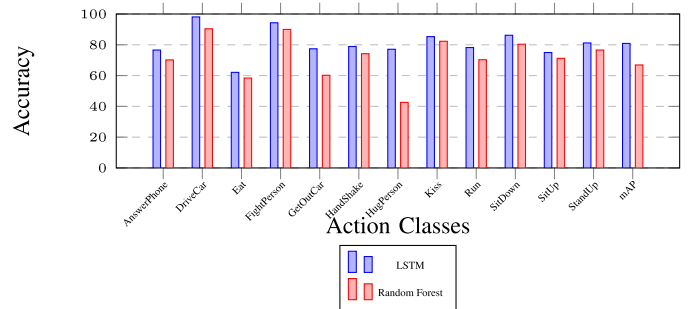


Fig. 10. mAP comparison for Random Forest Classifier and LSTM for Hollywood2 dataset. Motion:static ratio of 20:80 is used. mAP is significantly higher when the temporal dynamics of sub events are captured.

not capture sequential dynamic patterns, and compare it with the results obtained by the LSTM network. The results are shown in Fig. 9 and Fig. 10.

As it is evident from the results in in Fig. 9 and Fig. 10, LSTM network significantly outperforms the random forest classifier for both datasets. In Hollywood2, the LSTM network

wins by a 14% margin. In UCF-11, the LSTM network wins by a 12% margin. Therefore, it can be concluded that, exploiting temporal patterns of sub activities, benefits complex action classification.

V. CONCLUSION

This paper presents an end-to-end system for action classification which operates on both static and motion features. Our approach relies on deep features, for creating static vectors, and *motion tubes* for motion features. Motion tubes are a novel concept we introduce in this paper which can be used to track individual actors or objects across frames, and model micro level actions. We present three novel methods, based on Cholesky transformation, variance ratio, and PCA, for efficient combining of features from different domains, which is a vital requirement in action classification. Cholesky method provides the power to control the contribution of each domain in exact numbers, and variance ratio based method mathematically provides an optimum ratio for contribution. We show that these mathematical and experimental values agree with each other. We run experiments to show that the accuracy depends on the ratio of this contribution, and the optimum contribution of static and motion domains may vary depending on the richness of motion information. In short, our work demonstrates how the action classification accuracy varies with the combination ratio of static and motion features, while establishing the existence of an optimum combination ratio for a given test dataset.

In addition, we note that this variation of optimum ratio between static and motion feature contribution towards the final classification accuracy being dependent on the dataset demonstrates the possibility of this optimum ratio varying dynamically over time for even a given single video clip. Thus, based on the temporal variation of motion richness of any given video clip, the optimum contribution ratio has a possibility of varying dynamically. We hope to work on this in future.

Through our experiments we also show that our static and motion features are complementary, and contribute to the final result. We also compare our three fusion algorithms, and show that the Cholesky based method is superior, although all three of them give impressive results. We also model the temporal progression of sub-events using an LSTM network. Experimental results indicate that this is indeed beneficial, compared to using models which do not capture temporal dynamics. Comparison of our work with multiple state-of-the-art algorithms, on the popular datasets, UCF-11, and Hollywood2, show that our system performs better. The work on the HMDB51 dataset shows our system to be on par with the state-of-the-art.

In the future, it would be interesting to improve the motion tubes, so that, it can maintain an identity over each actor object. While it is mostly the case even in the present system, there is no guarantee. Also, in this work, the emphasis is on accelerating the per-actor micro action generation, initially we detect individual objects in first frame, and subsequently track those along motion tubes in the following frames. In the

future, exploring more powerful methods to describe micro actions inside motion tubes would be interesting, since it may increase the distinctiveness of the motion features and contribute well to the final accuracy.

REFERENCES

- [1] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Columbus, OH, USA, Jun. 2014, pp. 580–587.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, Lake Tahoe, NV, USA, Dec. 2012, pp. 1097–1105.
- [3] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu, "Action recognition by dense trajectories," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Colorado Springs, CO, USA, Jun. 2011, pp. 3169–3176.
- [4] H. Wang and C. Schmid, "Action recognition with improved trajectories," in *Proc. IEEE Int. Conf. Comput. Vis.*, Sydney, NSW, Australia, Dec. 2013, pp. 3551–3558.
- [5] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Proc. 28th Annu. Conf. Neural Inf. Process. Syst.*, Montreal, QC, Canada, Dec. 2014, pp. 568–576.
- [6] S. Ramasinghe and R. Rodrigo, "Action recognition by single stream convolutional neural networks: An approach using combined motion and static information," in *Proc. 3rd IAPR Asian Conf. Pattern Recognit. (ACPR)*, Kuala Lumpur, Malaysia, Nov. 2015, pp. 101–105.
- [7] R. Chaudhry, A. Ravichandran, G. Hager, and R. Vidal, "Histograms of oriented optical flow and Binet-Cauchy kernels on nonlinear dynamical systems for the recognition of human actions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Miami, FL, USA, Jun. 2009, pp. 1932–1939.
- [8] O. Russakovsky *et al.*, "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, Dec. 2015.
- [9] J. Liu, J. Luo, and M. Shah, "Recognizing realistic actions from videos 'in the wild,'" in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Miami, FL, USA, Jun. 2009, pp. 1996–2003.
- [10] M. Marszalek, I. Laptev, and C. Schmid, "Actions in context," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Miami, FL, USA, Jun. 2009, pp. 2929–2936.
- [11] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre, "HMDB: A large video database for human motion recognition," in *Proc. Int. Conf. Comput. Vis. (ICCV)*, Nov. 2011, pp. 2556–2563.
- [12] G. Gkioxari and J. Malik, "Finding action tubes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Boston, MA, USA, Jun. 2015, pp. 759–768.
- [13] M. Jain, J. van Gemert, H. Jégou, P. Bouthemy, and C. G. M. Snoek, "Action localization with tubelets from motion," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Columbus, OH, USA, Jun. 2014, pp. 740–747.
- [14] J. C. van Gemert, M. Jain, E. Gati, and C. G. M. Snoek, "APT: Action localization proposals from dense trajectories," in *Proc. Brit. Mach. Vis. Conf. (BMVC)*, vol. 2. Swansea, Wales, Sep. 2015, p. 4.
- [15] L. Wang, Y. Qiao, and X. Tang, "Action recognition with trajectory-pooled deep-convolutional descriptors," in *Proc. CVPR*, Boston, MA, USA, Jun. 2015, pp. 4305–4314.
- [16] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld, "Learning realistic human actions from movies," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Anchorage, AL, USA, Jun. 2008, pp. 1–8.
- [17] C. Schudt, I. Laptev, and B. Caputo, "Recognizing human actions: A local SVM approach," in *Proc. 17th Int. Conf. Pattern Recognit.*, vol. 3. Cambridge, U.K., Aug. 2004, pp. 32–36.
- [18] Y. Ke, R. Sukthankar, and M. Hebert, "Efficient visual event detection using volumetric features," in *Proc. 10th IEEE Int. Conf. Comput. Vis. (ICCV)*, vol. 1. Beijing, China, Jan. 2005, pp. 166–173.
- [19] E. Shechtman and M. Irani, "Space-time behavior based correlation," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, vol. 1. San Diego, CA, USA, Jun. 2005, pp. 405–412.
- [20] A. Klaser, M. Marszalek, and C. Schmid, "A spatio-temporal descriptor based on 3D-gradients," in *Proc. 19th Brit. Mach. Vis. Conf.*, Leeds, U.K., Sep. 2008, pp. 275:1–275:10.
- [21] T.-H. Yu, T.-K. Kim, and R. Cipolla, "Real-time action recognition by spatiotemporal semantic and structural forest," in *Proc. Brit. Mach. Vis. Conf.*, vol. 2. Aberystwyth, Wales, Aug. 2010, p. 6.

- [22] C. Feichtenhofer, A. Pinz and A. Zisserman, "Convolutional two-stream network fusion for video action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Las Vegas, NV, USA, 2016, pp. 1933–1941.
- [23] S. Ji, W. Xu, M. Yang, and K. Yu, "3D convolutional neural networks for human action recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 221–231, Jan. 2013.
- [24] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3D convolutional networks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Santiago, Chile, Dec. 2015, pp. 4489–4497.
- [25] H.-J. Kim, J. S. Lee, and H.-S. Yang, "Human action recognition using a modified convolutional neural network," in *Proc. Int. Symp. Neural Netw.*, Nanjing, China, Jun. 2007, pp. 715–723.
- [26] L. Wang, Y. Qiao, and X. Tang, "Action recognition and detection by combining motion and appearance features," in *Proc. ECCV*, Zürich, Switzerland, Sep. 2014.
- [27] B. Fernando, E. Gavves, M. J. Oramas, A. Ghodrati, and T. Tuytelaars, "Modeling video evolution for action recognition," in *Proc. 28th IEEE Conf. Comput. Vis. Pattern Recognit.*, Boston, MA, USA, Jun. 2015, pp. 5378–5387.
- [28] Y. Wang and G. Mori, "Hidden part models for human action recognition: Probabilistic versus max margin," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 7, pp. 1310–1323, Jul. 2011.
- [29] D. Wu and L. Shao, "Leveraging hierarchical parametric networks for skeletal joints based action segmentation and recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Columbus, OH, USA, Jun. 2014, pp. 724–731.
- [30] Y. Song, L.-P. Morency, and R. Davis, "Action recognition by hierarchical sequence summarization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Portland, OR, USA, Jun. 2013, pp. 3562–3569.
- [31] M. Rohrbach, M. Regneri, M. Andriluka, S. Amin, M. Pinkal, and B. Schiele, "Script data for attribute-based recognition of composite activities," in *Proc. Eur. Conf. Comput. Vis.*, Florence, Italy, Oct. 2012, pp. 144–157.
- [32] S. Bhattacharya, M. M. Kalayeh, R. Sukthankar, and M. Shah, "Recognition of complex events: Exploiting temporal dynamics between underlying concepts," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Columbus, OH, USA, Jun. 2014, pp. 2235–2242.
- [33] W. Li, Q. Yu, H. Sawhney, and N. Vasconcelos, "Recognizing activities via bag of words for attribute dynamics," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Portland, OR, USA, Jun. 2013, pp. 2587–2594.
- [34] E. A. Jackson, *Perspectives of Nonlinear Dynamics*, vol. 1. Cambridge, U.K.: CUP Archive, 1992.
- [35] T. Kailath, "A view of three decades of linear filtering theory," *IEEE Trans. Inf. Theory*, vol. 20, no. 2, pp. 146–181, Mar. 1974.
- [36] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [37] J. Y.-H. Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici, "Beyond short snippets: Deep networks for video classification," in *Proc. CVPR*, Boston, MA, USA, Jun. 2015, pp. 4694–4702.
- [38] J. Donahue *et al.*, "Long-term recurrent convolutional networks for visual recognition and description," in *Proc. CVPR*, Boston, MA, USA, Jun. 2015, pp. 2625–2634.
- [39] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proc. 7th Int. Joint Conf. Artif.*, 1981, pp. 674–679.
- [40] A. Graves, "Supervised sequence labelling with recurrent neural networks," Ph.D. dissertation, Faculty Comput. Sci., Tech. Univ. Munich, Munich, Germany, 2008.
- [41] E. Vig, M. Dorr, and D. Cox, "Space-variant descriptor sampling for action recognition based on saliency and eye movements," in *Proc. Eur. Conf. Comput. Vis.*, Florence, Italy, Oct. 2012, pp. 84–97.
- [42] N. Iizler-Cinbis and S. Sclaroff, "Object, scene and actions: Combining multiple features for human action recognition," in *Proc. Eur. Conf. Comput. Vis.*, Crete, Greece, Sep. 2010, pp. 494–507.
- [43] Y.-G. Jiang, Q. Dai, X. Xue, W. Liu, and C.-W. Ngo, "Trajectory-based modeling of human actions with motion reference points," in *Proc. Eur. Conf. Comput. Vis.*, Florence, Italy, Oct. 2012, pp. 425–438.
- [44] S. Mathe and C. Sminchisescu, "Dynamic eye movement datasets and learnt saliency models for visual action recognition," in *Computer Vision—ECCV*. Florence, Italy: Springer, Oct. 2012, pp. 842–856.
- [45] Y. Zhu and S. Newsam, "Depth2Action: Exploring embedded depth for large-scale action recognition," *CoRR*, 2016. [Online]. Available: <http://arxiv.org/abs/1608.04339>
- [46] M. Jain, H. Jégou, and P. Bouthemy, "Better exploiting motion for better action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Portland, OR, USA, Jun. 2013, pp. 2555–2562.
- [47] M. M. Ullah, S. N. Parizi, and I. Laptev, "Improving bag-of-features action recognition with non-local cues," in *Proc. Brit. Mach. Vis. Conf.*, vol. 10. Aberystwyth, Wales, Aug. 2010, pp. 1–95.



Sameera Ramasinghe received the B.Sc.Eng. degree in electronics and telecommunication from the University of Moratuwa, Sri Lanka, in 2014, where he is currently pursuing the M.Phil. degree. He is a co-founder and a research engineer with ConsientAI, a startup focused on AI technologies. His current research interests are machine learning and computer vision.



Jathushan Rajasegaran is currently pursuing the B.Sc.Eng. degree in electronics and telecommunication with the University of Moratuwa, Sri Lanka. His current research interests are machine learning, big data analysis, and data privacy.



Vinoj Jayasundara is currently pursuing the B.Sc.Eng. degree (Hons.) in electronics and telecommunication with the University of Moratuwa, Sri Lanka. His current research interests are machine learning, big data analytics, machine vision, and activity recognition.



Kanchana Ranasinghe is currently pursuing the B.Sc.Eng. degree in electronics and telecommunication with the University of Moratuwa, Sri Lanka. His current research interests are machine learning, computer vision, and pattern recognition.



Ranga Rodrigo received the B.Sc.Eng. degree (Hons.) from the University of Moratuwa, Moratuwa, Sri Lanka, in 2001, and the M.E.Sc. and Ph.D. degrees from Western University, London, ON, Canada, in 2004 and 2008, respectively. He has been with the Department of Electronic and Telecommunication Engineering, the University of Moratuwa, since 2008, where he is currently a Senior Lecturer. He is currently involved in feature tracking, reconstruction, and activity recognition. His research interests are in the general area of computer vision.



Ajith A. Pasqual received the B.Sc.Eng. degree (Hons.) in electronic and telecommunication engineering from the University of Moratuwa, Sri Lanka, in 1993, and the M.Eng. and Ph.D. degrees in computer vision from The University of Tokyo in 1998 and 2001, respectively. He was the former Head of Department of Electronic and Telecommunication Engineering, University of Moratuwa, where he is currently a Senior Lecturer. He leads the Reconfigurable Digital Systems Research Group, University of Moratuwa, which is currently involved in the area of hardware acceleration, novel architectures for application specific processors, and SoCs to improve performance and power efficiency. He is the Founder of the first Semiconductor Startup Company in Sri Lanka Paraqum Technologies, which is involved in developing high performance hardware decoder and encoder for the newest Video Compression Standard H.265/HEVC and network analytics equipment. His primary research interests are in application processors, machine vision, processor and SoC architectures.