

Activity Recognition based on a Magnitude-Orientation Stream Network

Carlos Caetano, Victor H. C. de Melo, Jefersson A. dos Santos, William Robson Schwartz
Smart Surveillance Interest Group, Department of Computer Science
Universidade Federal de Minas Gerais, Belo Horizonte, Brazil
{carlos.caetano,victorhcmelo,jefersson,william}@dcc.ufmg.br

Abstract—The temporal component of videos provides an important clue for activity recognition, as a number of activities can be reliably recognized based on the motion information. In view of that, this work proposes a novel temporal stream for two-stream convolutional networks based on images computed from the optical flow magnitude and orientation, named Magnitude-Orientation Stream (MOS), to learn the motion in a better and richer manner. Our method applies simple non-linear transformations on the vertical and horizontal components of the optical flow to generate input images for the temporal stream. Experimental results, carried on two well-known datasets (HMDB51 and UCF101), demonstrate that using our proposed temporal stream as input to existing neural network architectures can improve their performance for activity recognition. Results demonstrate that our temporal stream provides complementary information able to improve the classical two-stream methods, indicating the suitability of our approach to be used as a temporal video representation.

I. INTRODUCTION

Human activity recognition has been used in many real-world applications. In environments that require a higher level of security, surveillance systems can be used to detect and prevent abnormal or suspicious activities such as robberies and kidnappings. In addition, human activity recognition can be employed in systems for video retrieval, so that a user is able to search for videos containing specific activities. Another type of application is in health care, such as activities of daily living monitoring systems.

Surveillance applications have traditionally relied on network cameras monitored by human operators that must be aware of the activities carried out by people who are in the camera field of view. With the recent growth in the number of cameras to be analyzed, the efficiency and accuracy of human operators has reached the limit [1]. Therefore, security agencies have attempted computer vision-based solutions to replace or assist the human operator. Automatic recognition of suspicious activities is a problem that has attracted the attention of researchers in the area [2]–[5].

A significant portion of the progress on activity recognition task has been achieved with the design of discriminative feature descriptors exploring temporal information. Such information is based on motion analysis and is very important to represent the video in a more discriminative space, allowing the improvement of activity recognition.

Over the last decade, most of the works focused on designing handcrafted local feature descriptors [6]–[9] or on

encoding schemes using mid-level representations, such as Bag-of-Words (BoW) [10] or Fisher vector (FV) [11], followed by Support Vector Machines (SVM) classifier. Nowadays, large efforts have been directed to the employment of deep Convolutional Neural Networks (CNNs). These architectures learn hierarchical layers of representations to perform pattern recognition and have demonstrated impressive results on many pattern recognition tasks (e.g., image classification [12] and face recognition [13]). Although the excellent improvements achieved in such tasks, activity recognition lacks on performance when using CNNs. Many works [14]–[16] point that the potential reason behind such gap falls in two cases: (i) current datasets do not have enough videos for training and are too much noisy; and (ii) current CNN architectures are still not able to handle temporal information (or to take full advantage of it), consequently letting spatial (appearance) information prevail.

A major breakthrough spatiotemporal information representation was achieved by Simonyan and Zisserman [17], who directly incorporated motion information by using optical flow instead of learning it from scratch, showing significant improvement over other approaches. Known as two-stream network, their architecture is composed of two stream of data: (i) spatial network, which takes as input the raw RGB pixels; and (ii) temporal network, which takes as input dense optical flow displacement fields (vertical and horizontal components) computed across the frames. The final predictions are computed as the average of the output scores from the two streams.

To further improve the representation of spatiotemporal information, this work introduces a new temporal stream for the two-stream networks to perform activity recognition, named *Magnitude-Orientation Stream* (MOS). The method is based on non-linear transformations on the optical flow components to generate input images for the temporal stream. Our hypothesis is based on the assumption that the motion information on a video sequence can be described by the spatial relationship contained on the local neighborhood of magnitude and orientation extracted from the optical flow. More specifically, we assume that the motion information is adequately specified by fields of magnitude and orientation. In view of that, our method captures not only the displacement, by using orientation, but also magnitude providing information regarding the velocity of the movement.

In the literature, magnitude and orientation information

are often used to describe motion information in various local handcrafted-based features, such as Motion Boundary Histogram (MBH) [18], Histogram of Oriented Flow (HOF) [7], Histograms of Optical Flow Orientation and Magnitude (HOFM) [19] and Optical Flow Co-occurrence Matrices (OFCM) [20]. However, none of the aforementioned methods used such information on an end to end learning scheme with a CNN, as the proposed approach does.

According to the experimental results, our proposed temporal stream used as input to existing neural network architectures is able to recognize activities accurately on two well-know datasets (UCF101 [21] and HMDB51 [22]) outperforming the results achieved by the original two-stream network as well as other deep networks available in the literature.

II. RELATED WORKS

In this section, we present a literature review of works that are close to the idea proposed in our approach. These methods can be categorized by: (i) temporal information extracted from videos through the use of handcrafted local feature descriptors (Section II-A), and (ii) recent works that employ neural networks to learn temporal information (Section II-B).

A. Methods based on Handcrafted Feature Descriptors

To characterize motion and appearance of local features, Laptev et al. [7] computed histogram descriptors of space-time volumes in the neighborhood of detected points. Each volume is subdivided into a grid of cuboids and, for each cuboid, they compute Histogram of Oriented Gradients (HOG) [24] and Histogram of Optical Flow (HOF). The HOG descriptor is computed by dividing the cuboid into regions and accumulating a histogram binned by gradient directions over the pixels, while HOF is binned according to the flow orientations and weighted by according its magnitude. Then, normalized histograms are concatenated and named HOG-HOF.

Dalal et al. [18] introduced the Motion Boundary Histogram (MBH). First applied to human detection, the motion boundary coding scheme captures the local orientations of motion edges based on HOG feature descriptors [24]. Treating the horizontal and vertical components of the optical flow as independent “images”, the authors take their local gradients separately, find the corresponding magnitudes and orientations and use these as weighted votes to the local orientation histograms. Later on, the MBH was used on several works to describe motion information for activity recognition [7], [9], [25].

The HOG feature descriptor was extended by Kläser et al. [8], named as HOG3D. It is based on histograms of 3D gradient orientations computed using an integral video representation. The gradients are binned into regular polyhedrons in a multi-scale fashion in space and time. Therefore, HOG3D combines appearance and motion information.

Aiming at encoding both local static appearance and motion information, as in the HOG3D, but avoiding high dimensionality and a relatively expensive quantization cost, Shi et al. [26] proposed the Gradient Boundary Histograms (GBH). Instead of using image gradients, the authors use time-derivatives of

image gradients to emphasize moving edge boundaries. For each frame, they compute image gradients and apply temporal filtering over two consecutive gradient images. Then, they compute the magnitude and orientation for each pixel which are used to build a histogram of orientation as in HOG.

Colque et al. [19] developed a feature called Histograms of Optical Flow Orientation and Magnitude (HOFM). Different from HOF that only encodes orientation information, HOFM captures the orientation and the magnitude of flow vectors providing information regarding the velocity of the moving objects. They build a 3D matrix based on the orientation and magnitude information provided by the optical flow field, where each line corresponds to a given orientation range and each column to the magnitude ranges. The authors then extended it to capture information regarding appearance and density of regions by encoding the entropy of the orientation flow [27].

Aiming at capturing richer information from the optical flow, Caetano et al. [20] proposed the Optical Flow Co-occurrence Matrices (OFCM). The descriptor is based on the extraction of a set of statistical measures from co-occurrence matrices computed using the magnitude and orientation from optical flow information. Their hypothesis for designing the OFCM is based on the assumption that the motion information on a video sequence can be described by the spatial relationship contained on local neighborhoods of the flow field.

A major breakthrough on local feature-based approaches was achieved by Wang et al. [9] which proposed an method to describe videos by dense trajectories. Trajectory shapes encode local motion information by tracking spatial interest points over time. To generate the trajectories, they sample interest points in space and time, and track them based on displacement information using an efficient dense optical flow algorithm. The HOG, HOF and MBH feature descriptors are used to describe the trajectories which are then encoded by Bag-of-Words (BoW) mid-level representation. Afterwards, the authors improved it to the Improved Dense Trajectories (IDT) [25] using the homography between consecutive frames to estimate the camera motion and Fisher vector encoding.

Although there are many approaches based on local feature descriptors, these works often require over engineering (e.g., feature extraction, mid-level representation and classifier training). On contrary, CNNs are a class of deep learning models that replace all engineering with a single neural network trained end to end from pixel values to classifier outputs [28].

B. Methods based on Neural Network Approaches

Convolutional Neural Networks (CNNs) have achieved impressive state-of-the-art results on image classification [12]. Therefore, many works have tried to apply CNNs to learn spatiotemporal information for activity recognition task. A natural choice, the 3D convolutional network was presented by Ji et al. [29], where they tried to learn both appearance and motion features with 3D convolution operations. Their method works by stacking consecutive segments of human subjects in videos and by applying 3D convolutions over such volume

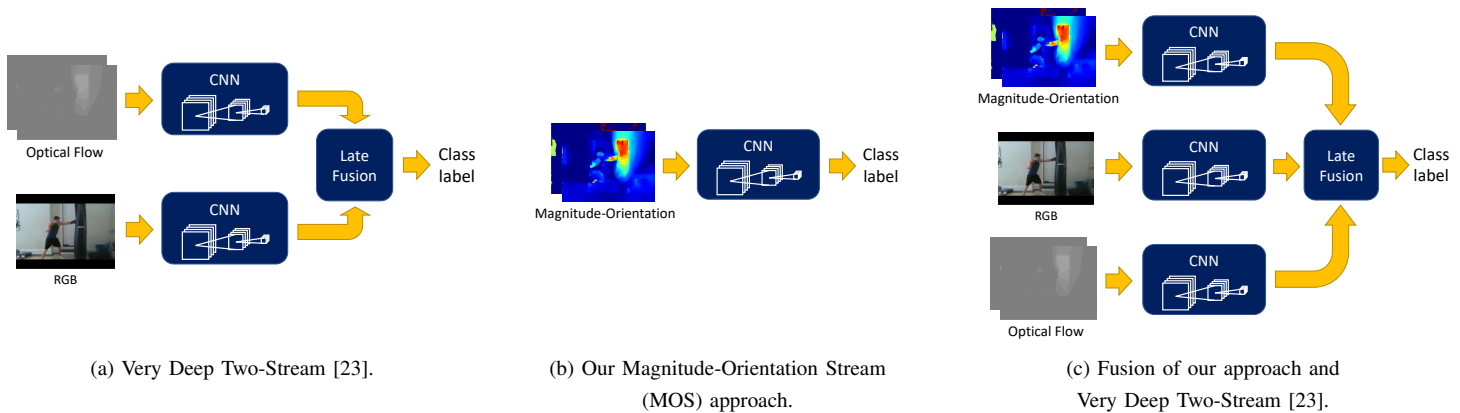


Fig. 1. Architectures considered in this work for extracting spatiotemporal information.

aiming that the first layer learns spatiotemporal features. Tran et al. [30] also explored 3D CNNs. However, in contrast with Ji et al. [29], their method takes full video frames as inputs and does not rely on any preprocessing.

Karpathy et al. [28] also used CNN aiming to learn motion features. The authors investigated different temporal information fusion schemes, learning local motion direction/speed with global information. Although significant gains in accuracy compared to the works based on handcrafted features, only little improvement was achieved when compared to single-frame CNN models, showing that the current CNN architectures are unable to efficiently learn motion features.

A major breakthrough was achieved by Simonyan and Zisserman [17]. Instead of trying to learn motion information as Karpathy et al. [28] and Tran et al. [30], the authors incorporated it by using optical flow. Known as two-stream network, their architecture is composed of two stream of data: (i) spatial network, which takes as input the raw RGB pixels; and (ii) temporal network, which takes as input dense optical flow displacements computed across the frames. Final predictions are computed as the average of the output scores from the two streams, showing significant improvement over other approaches. Our method differs from them by capturing not only the displacement but also velocity information provided by optical flow magnitude.

By employing the aforementioned two-stream network, Wang et al. [23] conducted experiments showing the impact on results when changing the network architecture. In addition, they also introduced some data augmentation techniques to improve the network training. To that end, the authors used three distinct architectures (ClarifaiNet [31], GoogLeNet [32] and VGG-16 [33]) showing that the best results are achieved by VGG-16 deeper architecture. Afterwards, the authors improved it to the Temporal Segment Networks (TSN) [34] by studying different types of input modalities to two-stream and by employing the Inception with batch normalization network architecture [35].

Perez et al. [36] used MPEG motion vectors [37] as a different input for a two-stream network to explore temporal information. Such vectors are used to perform motion esti-

mation in video compression where pixels are grouped in macroblocks and motion vectors are then computed for each block. They show that both optical flow and MPEG motion vectors provide equivalent accuracies, but the latter allows a more efficient implementation.

To make a spatial network learn to relate which parts of the image are moving, Park et al. [15] proposed a feature amplification technique by using magnitude information of the optical flow on the spatial network. To that end, they extract features maps of the last convolutional layer of the spatial network, compute optical flow magnitudes and resize it to be the same size of the previously extracted feature maps. Finally, they perform element-wise product to amplify the activations. Our work differs from them in that we use the magnitude information right on the beginning of the network, letting it learn how the velocity information contributes on the activity recognition process.

As it can be inferred from the reviewed methods, most of them use either convolution operations over raw pixels or optical flow to model temporal information. The former do not decouple spatial and temporal information, letting appearance information prevail [14], while the latter approaches rely on horizontal and vertical components of the optical flow. Despite the optical flow-based methods produce promising results, they focus only on displacement information. In view of that, aiming at capturing more information from the optical flow, our method captures not only the displacement, by using orientation, but also captures the magnitude providing information regarding the velocity of the movement.

III. PROPOSED APPROACH

In this section, we present our approach for performing activity recognition with our proposed *Magnitude-Orientation Stream* (MOS). For completeness, we first present the basic concepts of the Very Deep Two-Stream [23], which is the network architecture we use to learn the data representation based on the magnitude and orientation. Then, we detail our method showing how to incorporate magnitude and orientation as temporal information for the network input.

A. Very Deep Two-Stream

Motivated by the successful results achieved by deep architectures (e.g., VGG-16) in object recognition task, Wang et al. [23] improved the two-stream network by adapting it to use the VGG-16 on activity recognition, which they called Very Deep Two-Stream CNN. As mentioned on Section II, the two-stream network is composed by two different networks receiving distinct flows of data, spatial and temporal. The spatial stream receives as input the RGB frames while the temporal stream receives an optical flow image as input.

The spatial network is built on a single frame image and, therefore, its architecture is the same as those for object recognition on the image domain. Thus, at each iteration of the training step, 256 training videos are uniformly sampled across the classes and a single frame is randomly selected. Moreover, to avoid overfitting, the authors employ two data augmentation techniques: (i) cropping and flipping four corners and the center of the frame; and (ii) a multi-scale cropping method than randomly sampling the cropping width and height from 256, 224, 192, 168. Finally, they resize the cropped regions to $224 \times 224 \times 3$.

The temporal network receives images of optical flow as input. The process for computing the optical flow is explained as follows. For each frame F on time t , optical flow O_t is computed considering F_t and F_{t+1} . The resulting optical flow O_t is composed by two channels: (i) \mathcal{O}_t^x , denoting an image containing the x (horizontal) displacement field; and (ii) \mathcal{O}_t^y , denoting an image containing the y (vertical) displacement field. Moreover, to avoid storing the displacement fields as floats, the horizontal and vertical components of the flow are linearly rescaled to a $[0, 255]$ interval as

$$\mathcal{I}_{t,i,j}^f = \begin{cases} 0, & \text{if } \mathcal{O}_{t,i,j}^f < l \\ 255, & \text{if } \mathcal{O}_{t,i,j}^f > h \\ 255 \times \frac{(\mathcal{O}_{t,i,j}^f - l)}{(h-l)}, & \text{otherwise} \end{cases}, \quad (1)$$

where f represents the image channel (flow component x or y), h is the higher bound maximum optical flow value, l is the lower bound minimum optical flow value and \mathcal{I}^f the optical flow image. The same data augmentation techniques used in spatial network are used in the temporal stream. Finally, the input of the temporal network is composed by stacking 10 randomly images \mathcal{I}^f of optical flow fields ($224 \times 224 \times 20$) [17].

To perform the combination of the two networks, a late fusion scheme is employed by using a weighted linear combination of their prediction scores, where the weight is set as 2 for temporal network and 1 for spatial network, giving, therefore, more importance to the temporal information. Figure 1(a) illustrates the Deep Two-Stream network.

B. Magnitude-Orientation Stream

Our Magnitude-Orientation Stream (MOS) follows the same fundamentals as the Very Deep Two-Stream. However, aiming at extracting more information from the optical flow, MOS captures the displacement information by using orientation of the optical flow and the velocity of the movement considering

the optical flow magnitude. The spatial relationship contained on local neighborhoods of magnitude and orientation captures not only displacement, by using orientation, but also magnitude providing information regarding the velocity of the movement. The method is based on non-linear transformations on the optical flow components aiming to generate input images for the temporal stream. To incorporate such information on the temporal stream, we compute the dense optical flow as in [23]. In this way, for each video composed by n frames, we compute $n - 1$ optical flows \mathcal{O} . Once the optical flow is available, we compute the magnitude and orientation information as

$$M_{i,j} = \sqrt{(\mathcal{O}_{i,j}^x)^2 + (\mathcal{O}_{i,j}^y)^2} \quad (2)$$

and

$$\theta_{i,j} = \tan^{-1} \left(\frac{\mathcal{O}_{i,j}^y}{\mathcal{O}_{i,j}^x} \right), \quad (3)$$

where M and θ are the magnitude and orientation information, respectively.

Since the values obtained in M and θ are composed by real numbers, they are linearly rescaled to a $[0, 255]$ using Equation 1. Moreover, since the orientation values are estimated for every pixel of the optical flow, it can generate noisy values of regions of the image without any movement. Therefore, we performed a filtering on θ based on the values of M as

$$\theta'_{i,j} = \begin{cases} 0, & \text{if } M_{i,j} < m \\ \theta_{i,j}, & \text{otherwise} \end{cases}, \quad (4)$$

where m is a magnitude threshold value. Figure 2 illustrates a comparison between the magnitude and orientation information with the optical flow x and y displacements extracted from two consecutive frames.

With the rescaled magnitude and orientation information, which can be seen as two image channels, we use the same data augmentation techniques as in [23]. Therefore, the input is composed by 10 stacked images ($224 \times 224 \times 20$). Figure 1(b) illustrates the Magnitude-Orientation Stream network stages.

Finally, to incorporate spatial information to our approach, we employ a late fusion technique with the Very Deep Two-Stream network [23], as illustrated in Figure 1(c).

IV. EXPERIMENTAL RESULTS

This section describes the experimental results obtained with the proposed method for the activity recognition problem and performs comparisons to our baseline, the Very Deep Two-Stream network [23]. To isolate only the contribution brought by our method to the activity recognition problem, the baseline was tested on the same datasets with the same split of training and testing data. The evaluations are performed considering two well-known datasets for the activity recognition problem, the UCF101 [21] and the HMDB51 [22], in which we employ the evaluation protocols and metrics proposed by their authors.




Fig. 2. Comparison between optical flow displacement information and magnitude and orientation extracted from two consecutive frames (t and $t+1$) of an activity sample extracted from the UCF101 dataset [21].

A. Datasets

The *UCF101* [21] is an activity recognition dataset composed by videos collected from YouTube. It has a large diversity of activities and the presence of large variations in camera motion, object viewpoint, appearance, pose and scale, cluttered background, and illumination conditions. There are 13,320 videos from 101 activity categories grouped into 25 groups. Each group can consist of 4-7 videos of an activity. The videos from the same group may share some common features, such as similar background or similar viewpoint. We follow the original protocol using three train-test splits. The performance is evaluated by computing the average recognition across all classes over the three splits as in [23].

The *HMDB51* [22] is a realistic and challenging activity dataset composed of video clips from movies, the Prelinger archive, Internet, Youtube and Google videos, and comprised of 51 activity categories. It consists of 6,766 activity samples with a resolution of 240 pixels in height with preserved aspect ratio. We follow the original protocol using three train-test splits. The performance is evaluated by computing the average accuracy across all classes over the three splits.

B. Implementation Details

1) *Network Architecture*: We employed the VGG-16 deep convolutional network [33]. It is a convolutional architecture composed by 13 convolutional and three fully-connected layers with smaller convolutional size, stride and pooling window (3×3 , 1×1 and 2×2 , respectively). We opted for using such

architecture since it demonstrated the best performance in [23]. In addition, we used a model compatible with a modified version of the Caffe framework [38] made available¹.

2) *Pre-training*: As stated by [23], the UCF101 dataset training split is very small to train a deep convolutional network (the same applies to HMDB51, since it has less videos). In view of that, a possible solution used by several works [14], [17], [23], [34] is to use ImageNet models as the initialization for network training. In this way, here we also employed the ImageNet model as pre-training

3) *Training*: Following the implementation details used by our baseline [23], we set the learning rate initially to 0.005 and then decreases it at every 5,000 iterations dividing it by 10. The maximum iteration was set as 15,000. We kept the same schedule for all training sets.

Similarly to [17], [23], the network weights are learned using the mini-batch stochastic gradient descent with a momentum set to 0.9 and weight decay of 0.0005. We also set high dropout ratio for the fully connected layers (0.9 and 0.8).

Krizhevsky et al. [12] demonstrated that data augmentation techniques can be very effective to avoid overfitting. In view of that, we cropped and flipped four corners and the center of the frame. In addition, we applied a multi-scale cropping method and randomly sampled the cropping width and height from $\{256, 224, 192, 168\}$ (finally, we resize the cropped regions to 224×224). It is important to state that our baseline [23] used the same data augmentation procedure.

4) *Test*: To perform a fair comparison, we applied the same test scheme used by our baseline [23], described as follows. First, we sample 25 magnitude/orientation flow images for the testing. Then, from each of these, we obtain 10 convolutional network inputs (by cropping and flipping four corners and the center). Finally, the prediction score for the input video is obtained by averaging the sampled images scores and their crops. The same testing scheme was used by the original two-stream convolutional network [17]. For the fusion of MOS and other streams, we use a non-weighted linear combination of their prediction scores.

5) *Optical Flow Extraction*: As mentioned on Section III, the magnitude/orientation images are computed from the optical flow information. To that end, we extract the optical flow information using the TVL1 algorithm [39], implemented in OpenCV with CUDA. For the sake of comparison, our baseline [23] used the same optical flow algorithm. To obtain the magnitude and orientation images information we empirically set the parameters $h = 15$ and $l = -15$ to compute M ; and $h = 180$, $l = -180$ and $m = 128$ to compute θ' .

C. Evaluation

We report the activity recognition performance of our Magnitude-Orientation Stream in contrast with the baseline in Table I showing a comparison of our method to the three different streams of our baselines (Very Deep Spatial Stream, Very Deep Temporal Stream and Very Deep Two-Stream).

¹<https://github.com/yjxiang/caffe>

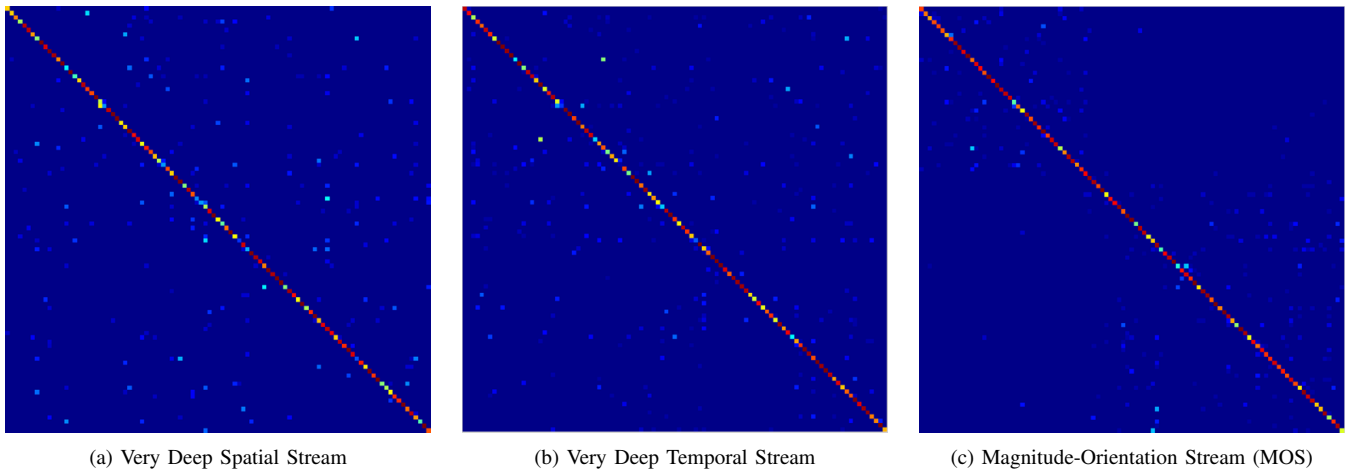


Fig. 3. Confusion matrices on UCF101 split 1. False positives and false negatives were highlighted to show where each method fails.

TABLE I

ACTIVITY RECOGNITION ACCURACY (%) RESULTS OF MAGNITUDE-ORIENTATION STREAM AND THE BASELINE ON UCF101 [21] ACTIVITY DATASET. RESULTS FOR THE BASELINE WERE OBTAINED RUNNING THE CODE PROVIDED BY THE AUTHORS [23]. NOTE THAT OUR RESULTS WERE ACHIEVED WITH ONLY OUR SINGLE MAGNITUDE-ORIENTATION STREAM (TEMPORAL INFORMATION) WHILE THE RESULTS OF [23] CONSIDER TWO STREAMS (SPATIAL AND TEMPORAL INFORMATION).

Approach		Split 1	Split 2	Split 3	Average
		Acc. (%)	Acc. (%)	Acc. (%)	Acc. (%)
Baseline	Very Deep Spatial Stream [23]	79.8	77.3	77.8	78.4
	Very Deep Temporal Stream [23]	85.7	88.2	87.4	87.0
	Very Deep Two-Stream [23]	90.9	91.6	91.6	91.4
Our results	Magnitude-Orientation Stream (MOS)	90.8	89.3	91.5	90.5
	Magnitude-Orientation Stream (MOS) + Very Deep Spatial Stream	93.1	91.9	92.6	92.5
	Magnitude-Orientation Stream (MOS) + Very Deep Temporal Stream	91.4	92.2	93.6	92.4
	Magnitude-Orientation Stream (MOS) + Very Deep Two-Stream	93.7	93.1	94.8	93.8

A considerable improvement was obtained with Magnitude-Orientation Stream when compared to the baseline single streams, reaching 90.8% of accuracy on split 1 of the UCF101 dataset. There is an improvement of 5.1 percentage points (p.p.) when compared to the Very Deep Temporal Stream [23] and 11.0 p.p. when compared to the Very Deep Spatial Stream [23]. Furthermore, it is worth noting that our best result using Magnitude-Orientation Stream on split 1 is close to the best one reported (Very Deep Two-Stream) which is obtained by using a combination of two different streams (spatial and temporal informations), while we only used our single Magnitude-Orientation Stream (temporal information). The same observations can be considered when analyzing the results of our temporal stream on splits 2 and 3. Therefore, such results can be considered remarkably good and confirm the advantages introduced by our approach.

Figure 3 shows the confusion matrices of Very Deep Spatial Stream, Very Deep Temporal Stream and our Magnitude-Orientation Stream for the UCF101 split 1 (we highlighted the false positives and false negatives to make it more visible on where each method fails). We can observe that our approach fails on classes that are more semantically closer to each

other², whereas the Very Deep Spatial Stream and the Very Deep Temporal Stream fails in a random manner. In addition, the three methods produce false positives and false negatives different from each other, indicating the possibility of fusion.

To exploit a possible complementarity of the three approaches (very deep spatial stream, very deep temporal stream and our magnitude-orientation stream), we combined the different streams by employing a late fusion technique using a weighted linear combination of their prediction scores. According to the results showed in Table I, any type of combination performed with our Magnitude-Orientation Stream provides better results than Very Deep Two-Stream, with the best result achieving an improvement of 2.4 p.p. over Very Deep Two-Stream.

Table II presents results on UCF101 dataset for many works. The first part of the table shows results of methods that extract temporal information using handcrafted features. We compare our MOS approach with the results of local feature-based methods, such as Bag-of-Words (BoW) + features, Fisher

²Since the actions on the confusion matrices are sorted according to its labels (e.g., ApplyEyeMakeup, ApplyLipstick, or BaseballPitch, Basketball, BasketballDunk), near regions in the confusion matrix denote semantically closer activities.

TABLE II

ACTIVITY RECOGNITION ACCURACY (%) RESULTS OF HANDCRAFTED METHODS, NEURAL NETWORKS (NN) METHODS AND OUR MAGNITUDE-ORIENTATION STREAM ON UCF101 ACTIVITY DATASET [21]. RESULTS FOR FEATURES + BoW WERE OBTAINED FROM [26] AND FEATURES + FV WERE OBTAINED FROM [40].

	Approach	UCF101 Acc. (%)
Handcrafted Methods	HOF + BoW [7]	61.8
	HOG-HOF + BoW [7]	71.8
	MBH + BoW [18]	77.1
	GBH + BoW [26]	68.5
	HOG3D + BoW [8]	61.4
	HOF + FV [7]	65.9
	HOG-HOF + FV [7]	75.4
	MBH + FV [18]	81.0
	GBH + FV [26]	74.2
	HOG3D + FV [8]	64.7
	IDT [25]	85.9
	IDT + higher FV [41]	87.9
	IDT + MVSV [42]	83.5
NN Methods	Deep Networks [28]	65.4
	Composite LSTM [43]	75.8
	C3D [30]	85.2
	Factorized CNN [44]	88.1
	Two-Stream [17]	88.0
	Two-Stream F [14]	92.5
	KVMF [45]	93.1
TSN [34]	94.2	
Our results	MOS	90.5
	MOS + Very Deep Spat. Stream	92.5
	MOS + Very Deep Temp. Stream	92.4
	MOS + Very Deep Two-Stream	93.8

vector (FV) + features, and Improved Dense Trajectories (IDT). The best result by such type of methods was achieved with IDT + higher FV [41], reaching 87.9%. Our best result using the proposed approach combined with Very Deep Two-Stream outperforms that by 5.9 percentage points.

The second part of Table II shows results achieved with neural networks (NN) approaches. According to the results, just using our Magnitude-Orientation Stream (MOS), we outperform many methods ([17], [28], [30], [43], [44]). In comparison with C3D [30], we outperform them by 5.3 p.p. using our temporal stream and 8.6 p.p. when combining it with Very Deep Two-Stream. This indicates that our magnitude orientation approach learns temporal information better than the approaches that perform 3D convolution operations directly. Moreover, such finding is very important since 3D convolutional operations are more computationally expensive than the 2D convolutional operations used in our approach. It is worth mentioning that we also improved the results achieved by the original two-stream [17] by 2.5 p.p. using our temporal stream and by 5.8 p.p. combining it with Very Deep Two-Stream. Finally, we can observe that our best result only did not outperform Wang et al. TSN method [34], however it presents results very close to it.

Finally, Table III presents the results achieved on HMDB51

TABLE III

ACTIVITY RECOGNITION ACCURACY (%) RESULTS OF HANDCRAFTED METHODS, NEURAL NETWORKS (NN) METHODS AND OUR MAGNITUDE-ORIENTATION STREAM ON HMDB51 ACTIVITY DATASET [22]. RESULTS FOR FEATURES + BoW WERE OBTAINED FROM [26] AND FEATURES + FV WERE OBTAINED FROM [40].

	Approach	HMDB51 Acc. (%)
Handcrafted Methods	HOF + BoW [7]	35.5
	HOG-HOF + BoW [7]	43.6
	MBH + BoW [18]	51.5
	GBH + BoW [26]	38.8
	HOG3D + BoW [8]	36.2
	OFCM + BoW [20]	56.9
	HOF + FV [7]	39.9
	HOG-HOF + FV [7]	45.6
	MBH + FV [18]	54.7
	GBH + FV [26]	44.7
	HOG3D + FV [8]	38.2
	IDT [25]	57.2
	IDT + higher FV [41]	61.1
	IDT + MVSV [42]	55.9
	NN Methods	Composite LSTM [43]
Factorized CNN [44]		59.1
Two-Stream [17]		59.4
Two-Stream F [14]		65.4
KVMF [45]		63.3
TSN [34]	69.4	
Our results	MOS	66.2

dataset with works of the literature that tried to handle temporal information by handcrafted features (first part) as well as recent works that employ neural networks (second part). Again, the best handcrafted method result is achieved by IDT + higher FV [41] reaching 61.1%. Our MOS approach outperformed it by 5.1 p.p. When compared with neural network methods, we were able to outperform many methods, e.g., [17], [28], [30], [43]–[45], using the proposed Magnitude-Orientation Stream (MOS). Furthermore, we were able to outperform the original two-stream [17] by 6.6 p.p. just using our temporal stream. Once more, our approach only did not outperform TSN [34]. Here, we believe this is because HMDB51 is a smaller dataset lacking training data³.

V. CONCLUSIONS AND FUTURE WORKS

In this work, we proposed a novel temporal stream for two-stream convolutional networks, named Magnitude-Orientation Stream (MOS). The method is based on simple non-linear transformations on the optical flow components generating input images composed of magnitude and orientation information. The spatial relationship contained on local neighborhoods of magnitude and orientation captures not only displacement, by using orientation, but also magnitude providing information regarding the velocity of the movement. We demonstrated

³The same statement was made by Simonyan and Zisserman [17], which evaluated different options for increasing the effective training set size of HMDB51.

that MOS outperforms all classic approaches based on local handcrafted features of the literature. Furthermore, simply by using only our temporal stream, we outperform original CNN two-stream approaches based on temporal and spatial information as well as other recent works that employ neural networks, suggesting its suitability to learn temporal information. Another interesting finding is that the combination of our temporal stream with the Very Deep Two-Stream method improves the activity recognition.

Directions to future works include the evaluation of the magnitude-orientation information with other distinct architectures, such as the Inception with batch normalization. Moreover, we intend to evaluate its behavior on other video-related problems, such as action detection or scene description.

ACKNOWLEDGMENTS

The authors would like to thank the Brazilian National Research Council – CNPq, the Minas Gerais Research Foundation – FAPEMIG (Grants APQ-00567-14 and PPM-00540-17) and the Coordination for the Improvement of Higher Education Personnel – CAPES (DeepEyes Project). The authors gratefully acknowledge the support of NVIDIA Corporation with the donation of the GeForce Titan X GPU used for this research.

REFERENCES

- [1] H. Keval, "CCTV Control Room Collaboration and Communication: Does it Work?" in *Human Centred Technology Workshop*, 2006.
- [2] S. Danafar and N. Gheissari, "Action recognition for surveillance applications using optic flow and SVM," in *ACCV*, 2007.
- [3] V. Reddy, C. Sanderson, and B. Lovell, "Improved anomaly detection in crowded scenes via cell-based analysis of foreground speed, size and texture," in *CVPRW*, 2011.
- [4] A. Wiliem, V. Madasu, W. Boles, and P. Yarlagadda, "A suspicious behaviour detection using a context space model for smart surveillance systems," *Comput. Vis. Image Underst.*, 2012.
- [5] J. Wang and Z. Xu, "Spatio-temporal texture modelling for real-time crowd anomaly detection," *Computer Vision and Image Understanding*, 2016.
- [6] P. Scovanner, S. Ali, and M. Shah, "A 3-dimensional Sift Descriptor and Its Application to Action Recognition," in *ACM MM*, 2007.
- [7] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld, "Learning realistic human actions from movies," in *CVPR*, 2008.
- [8] A. Kläser, M. Marszalek, and C. Schmid, "A Spatio-Temporal Descriptor Based on 3D-Gradients," in *BMVC*, 2008.
- [9] H. Wang, A. Klaser, C. Schmid, and C.-L. Liu, "Action recognition by dense trajectories," in *CVPR*, 2011.
- [10] J. Sivic and A. Zisserman, "Video Google: A Text Retrieval Approach to Object Matching in Videos," in *ICCV*, 2003.
- [11] J. Sánchez, F. Perronnin, T. Mensink, and J. Verbeek, "Image Classification with the Fisher Vector: Theory and Practice," *Int. J. Comput. Vision*, 2013.
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *NIPS*, 2012.
- [13] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A Unified Embedding for Face Recognition and Clustering," in *CVPR*, 2015.
- [14] C. Feichtenhofer, A. Pinz, and A. Zisserman, "Convolutional two-stream network fusion for video action recognition," in *CVPR*, 2016.
- [15] E. Park, X. Han, T. L. Berg, and A. C. Berg, "Combining multiple sources of knowledge in deep CNNs for action recognition," in *WACV*, 2016.
- [16] A. Diba, A. M. Pazandeh, and L. Van Gool, "Efficient Two-Stream Motion and Appearance 3D CNNs for Video Classification," in *ECCV*, 2016.
- [17] K. Simonyan and A. Zisserman, "Two-stream Convolutional Networks for Action Recognition in Videos," in *NIPS*, 2014.
- [18] N. Dalal, B. Triggs, and C. Schmid, "Human Detection Using Oriented Histograms of Flow and Appearance," in *ECCV*, 2006.
- [19] R. V. H. M. Colque, C. Caetano, and W. R. Schwartz, "Histograms of Optical Flow Orientation and Magnitude to Detect Anomalous Events in Videos," in *SIBGRAPI*, 2015.
- [20] C. Caetano, J. A. dos Santos, and W. R. Schwartz, "Optical Flow Co-occurrence Matrices: A novel spatiotemporal feature descriptor," in *ICPR*, 2016.
- [21] K. Soomroand, A. R. Zamir, and M. Shah, "UCF101: A Dataset of 101 Human Action Classes From Videos in The Wild," *CRCV-TR*, Tech. Rep., 2012.
- [22] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre, "HMDB: A large video database for human motion recognition," in *ICCV*, 2011.
- [23] L. Wang, Y. Xiong, Z. Wang, and Y. Qiao, "Towards Good Practices for Very Deep Two-Stream ConvNets," *CoRR*, 2015.
- [24] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *CVPR*, 2005.
- [25] H. Wang and C. Schmid, "Action Recognition with Improved Trajectories," in *ICCV*, 2013.
- [26] F. Shi, R. Laganiere, and E. Petriu, "Gradient Boundary Histograms for Action Recognition," in *WACV*, 2015.
- [27] R. V. H. M. Colque, C. Caetano, M. T. L. de Andrade, and W. R. Schwartz, "Histograms of Optical Flow Orientation and Magnitude and Entropy to Detect Anomalous Events in Videos," *IEEE Transactions on Circuits and Systems for Video Technology*, 2017.
- [28] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, "Large-Scale Video Classification with Convolutional Neural Networks," in *CVPR*, 2014.
- [29] S. Ji, W. Xu, M. Yang, and K. Yu, "3D Convolutional Neural Networks for Human Action Recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2013.
- [30] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning Spatiotemporal Features With 3D Convolutional Networks," in *ICCV*, 2015.
- [31] M. D. Zeiler and R. Fergus, *Visualizing and Understanding Convolutional Networks*. Springer International Publishing, 2014, pp. 818–833.
- [32] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going Deeper With Convolutions," in *CVPR*, 2015.
- [33] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in *ICLR*, 2015.
- [34] 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," in *ECCV*, 2016.
- [35] S. Ioffe and C. Szegedy, "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," in *ICML*, 2015.
- [36] M. Perez, S. Avila, D. Moreira, D. Moraes, V. Testoni, E. Valle, S. Goldenstein, and A. Rocha, "Video pornography detection through deep learning techniques and motion information," *Neurocomputing*, 2017.
- [37] I. E. G. Richardson, *H.264 and MPEG-4 Video Compression: Video Coding for Next-generation Multimedia*. John Wiley & Sons, Inc., 2003.
- [38] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional Architecture for Fast Feature Embedding," in *ACM MM*, 2014.
- [39] C. Zach, T. Pock, and H. Bischof, "A Duality Based Approach for Realtime TV-L1 Optical Flow," in *Proceedings of the 29th DAGM Conference on Pattern Recognition*, 2007.
- [40] F. Shi, "Local Part Model for Action Recognition in Realistic Videos," Ph.D. dissertation, School of Electrical Engineering and Computer Science, Faculty of Engineering, University of Ottawa, 2014.
- [41] X. Peng, L. Wang, X. Wang, and Y. Qiao, "Bag of visual words and fusion methods for action recognition: Comprehensive study and good practice," *Computer Vision and Image Understanding*, 2016.
- [42] Z. Cai, L. Wang, X. Peng, and Y. Qiao, "Multi-view Super Vector for Action Recognition," in *CVPR*, 2014.
- [43] N. Srivastava, E. Mansimov, and R. Salakhutdinov, "Unsupervised Learning of Video Representations Using LSTMs," in *ICML*, 2015.
- [44] L. Sun, K. Jia, D. Y. Yeung, and B. E. Shi, "Human Action Recognition Using Factorized Spatio-Temporal Convolutional Networks," in *ICCV*, 2015.
- [45] W. Zhu, J. Hu, G. Sun, X. Cao, and Y. Qiao, "A Key Volume Mining Deep Framework for Action Recognition," in *CVPR*, 2016.