

2018 24th International Conference on Pattern Recognition (ICPR). Piscataway, U.S. : Institute of Electrical and Electronics Engineers, Inc. pp. 566-571. ISSN 1051-4651 ISBN 9781538637883.

<https://dx.doi.org/10.1109/ICPR.2018.8545723>

Superframes, A Temporal Video Segmentation

Hajar Sadeghi Sokeh*, Vasileios Argyriou*, Dorothy Monekosso†, Paolo Remagnino*

*The Robot Vision Team, Faculty of Science, Engineering, and Computing
Kingston University, UK, London

{h.sadeghisokeh, vasileios.argyriou, p.remagnino}@kingston.ac.uk

†The Robot Vision Team, Faculty of Science, Engineering, and Computing, Leeds Beckett University, UK, Leeds
{d.n.monekosso}@leedsbeckett.ac.uk

Abstract—The goal of video segmentation is to turn video data into a set of concrete motion clusters that can be easily interpreted as building blocks of the video. There are some works on similar topics like detecting scene cuts in a video, but there is few specific research on clustering video data into the desired number of compact segments. It would be more intuitive, and more efficient, to work with perceptually meaningful entity obtained from a low-level grouping process which we call it ‘superframe’. This paper presents a new simple and efficient technique to detect superframes of similar content patterns in videos. We calculate the similarity of content-motion to obtain the strength of change between consecutive frames. With the help of existing optical flow technique using deep models, the proposed method is able to perform more accurate motion estimation efficiently. We also propose two criteria for measuring and comparing the performance of different algorithms on various databases. Experimental results on the videos from benchmark databases has demonstrated the effectiveness of the proposed method.

I. INTRODUCTION

In computer vision, many existing algorithms on video analysis use fixed number of frames for processing. For example optical flow or motion estimation techniques [1] and human activity recognition [2], [3]. However, it would be more intuitive, and more efficient, to work with perceptually meaningful entity obtained from a low-level grouping process which we call it ‘superframe’.

Similar to superpixels [4] which are key building blocks of many algorithms and significantly reduce the number of image primitives compared to pixels, superframes also do the same in time domain. They can be used in many different applications such as video segmentation [5], video summarization [6], and video saliency detection [7]. They are also useful in the design of a video database management system [8] that manages a collection of video data and provides content-based access to users [9]. Video data modeling, insertion, storage organization and management, and video data retrieval are among the basic problems that are addressed in a video database management system which can be solved more efficiently by using superframes. By temporal clustering of the video, it is easier to identify the significant segment of the video to achieve better representation, indexing, storage, and retrieval of the video data.

The main goal of this work is an automatic temporal clustering of a video by analyzing the visual content of the video and partitioning it into a set of units called superframes. This process can also be referred to as video data segmentation.

Each segment is defined as a continuous sequence of video frames which have no significant inter-frame difference in terms of their motion contents. Motion is the main criteria we use in this work, therefore we assume all the videos are taken from a single fixed camera.

There is little literature on specifically temporal segmenting of video, and some works on related areas. In this section, we briefly discuss the most relevant techniques to this work: temporal superpixels, scene cut, and video segmentation.

The main idea of using superpixels as primitives in image processing was introduced by Ren and Malik in [10]. Using superpixels instead of raw pixel data is even beneficial for video applications. Although until recently, superpixel algorithms were mainly on the still images, researchers started to apply them to the video sequences. There are some recent works on using the temporal connection between consecutive frames. Reso et al. [11] proposed a new method for generating superpixels in a video with temporal consistency. Their approach performs an energy-minimizing clustering using a hybrid clustering strategy for a multi-dimensional feature space. This space is separated into a global color subspace and multiple local spatial subspaces. A sliding window consisting multiple consecutive frames is used which is suitable for processing arbitrarily long video sequences.

A generative probabilistic model is proposed by Chang et al. in [12] for temporally consistent superpixels in video sequences which uses past and current frames and scales linearly with video length. They have presented a low-level video representation which is very related to a volumetric voxel [13], but still different in such a way that temporal superpixels are mainly designed for video data, whereas supervoxels are for 3D volumetric data.

Lee et al. in [14] developed a temporal superpixel algorithm based on proximity-weighted patch matching. They estimated superpixel motion vectors by combining the patch matching distances of neighboring superpixels and the target superpixel. Then, they initialized the superpixel label of each pixel in a frame, by mapping the superpixel labels in the previous frames using the motion vectors. Next, they refined the initial superpixels by updating the superpixel labels of boundary pixels iteratively based on a cost function. Finally, they performed some postprocessing including superpixel splitting, merging, and relabeling.

In video indexing, archiving and video communication such




Fig. 1. An example of superframes in a video from MAD database

as rate control, scene change detection plays an important role. It can be very challenging when scene changes are very small and sometimes other changes like brightness variation may cause false change detection. Many various methods have been proposed for scene change detection. Yi and Ling in [15], proposed a simple technique to detect sudden and unexpected scene change based on only pixel values without any motion estimation. They first screen out many non-scene change frames and then normalize the rest of the frames using a histogram equalization process.

A fully convolutional neural network has been used for shot boundary detection task in [16] by Gygli. He considered this work as a binary classification problem to correctly predict if a frame is part of the same shot as the previous frame or not. He also created a new dataset of synthetic data with one million frames to train this network.

Video segmentation aims to group perceptually and visually similar video frames into spatio-temporal regions, a method applicable to many higher-level tasks in computer vision such as activity recognition, object tracking, content-based retrieval, and visual enhancement. In [5], Grundmann et al. presented a technique for spatio-temporal segmentation of long video sequences. Their work is a generalization of Felzenszwalb and Huttenlocher’s [17] graph-based image segmentation technique. They use a hierarchical graph-based algorithm to make an initial over-segmentation of the video volume into relatively small space-time regions. They use optical flow as a region descriptor for graph nodes.

Kotsia et al. in [18] proposed using the 3D gradient correlation function operating at the frequency domain for action spotting in a video sequence. They used the 3D Fourier transform which is invariant to spatiotemporal changes and frame recording. In this work, the estimation of motion relies on the detection of the maximum of the cross-correlation function between two blocks of video frames [19].

Similar to all these tasks, in this work we propose a simple and efficient motion-based method to segment a video over time into compact episodes. This grouping leads to an increased computational efficiency for subsequent processing steps and allows for more complex algorithms computationally

infeasible on frame level. Our algorithm detects major changes in video streams, such as when an action begins and ends.

The only input to our method is a video and a number which shows the desired number of clusters in that video. The output will be the frame numbers which shows the boundaries between the uniform clusters. Visual features like colors and textures per frame are of low importance compare to motion features. This motivates us to employ motion information to detect big changes over the video frames. We use FlowNet-2 [1], a very accurate optical flow estimation with deep networks, to extract the motion between every subsequent frame. We then use both the average and the histogram of flow over video frames and compare our results over a baseline.

II. THE PROPOSED SUPERFRAME TECHNIQUE

Our superframe segmentation algorithm detects the boundary between temporal clusters in video frames. The superframe algorithm takes the number of desired clusters, K , as input, and generates superframes based on the motion similarity and proximity in the video frames. In other words, the histogram of magnitude is used with the direction of motion per frame and the frame position in the video as features to cluster video frames. Our superframe technique is motivated by SLIC [4], a superpixel algorithm for images, which we generalize it for video data (see Figure 1).

A. The proposed algorithm

Our model segments the video using motion cues. We assume the videos are taken with a fixed camera. Therefore, one can represent the type of motion of the foreground object by computing features from the optical flow. We first apply FlowNet-2 [1] to get the flow per video frame, and then we make a histogram of magnitude (HOM) and direction (HOD) of flow per frame to initialize K cluster centers C_1, \dots, C_K . Therefore each cluster center has 20 feature values as $C_k = [HOM_{1..11}, HOD_{1..8}, fr]^T$ with $k = [1, K]$. $HOM_{1..11}$ indicates ‘Histogram of Magnitude’ and $HOD_{1..8}$ indicates ‘Histogram of Direction’ for 8 different directions, and fr stands for the frame index in the video.

For a video with N frames, in the initialization step, there are K equally-sized superframes with approximately N/K frames. Since initially, the length of each superframe is $S = N/K$, like [4], we safely assume that the search area to find the best place for a cluster center is a $2S \times 2S$ area around each cluster center over the location of video frames.

Following the cluster center initialization, a distance measure is considered to specify each frame belongs to which cluster. We use D_s as a distance measure defined as follows:

Distance between cluster k and frame i is calculated by:

$$\begin{aligned} d_c &= \sqrt{\sum (X_k - X_i)^2}, \quad X = (x_1 \dots x_f) \\ d_s &= fr_k - fr_i \\ D_s &= \sqrt{\left(\frac{d_c}{m}\right)^2 + \left(\frac{d_s}{S}\right)^2} \end{aligned} \quad (1)$$

where x is a feature value and X is a feature vector of 19 values per video frame including 11 values for the histogram of the magnitude of flow and 8 values for the histogram of the direction of flow. We consider frame location separately as fr . S is the interval and m is a measure of compactness of a superframe which regarding the experiment we choose it as 10% of the input number of clusters in this work i.e. $0.1 * K$ to make the result comparison easier.

After the initialization of the K cluster centers, we then move each of them to the lowest gradient position in a neighborhood of 3 frames. The neighbourhood of 3 frames is chosen arbitrarily but reasonable. This is done to avoid choosing noisy frames. The gradient for frame i is computed as Eq. 2:

$$G(i) = \sqrt{\sum (X(i+1) - X(i-1))^2} \quad (2)$$

We associate each frame in the video with the nearest cluster center in the search area of $2S \times 2S$. When all frames are associated with the nearest cluster center, a new cluster center is computed as the average of flow values of all the frames belonging to that cluster. We repeat this process until convergence when the error is less than a threshold.

At the end of this process, we may have few clusters which their length is very short. So, we do a postprocessing step to merge these very small clusters to the closer left or right cluster. The whole algorithm is summarized in Algorithm 1

B. Evaluation criteria for the performance of the algorithm

Supapixel algorithms are usually assessed using two error metrics for evaluation of segmentation quality: boundary recall and under-segmentation error. Boundary recall measures how good a superframe segmentation adhere to the ground-truth boundaries. Therefore, higher boundary recall describes better adherence to video segment boundaries. Suppose $S = \{S_1, \dots, S_H\}$ be a superframe segmentation with H number of clusters and $G = \{G_1, \dots, G_L\}$ be a ground-truth segmentation with L number of clusters. Then, boundary recall $Rec(S, G)$ is defined as follows:

$$Rec(S, G) = \frac{TP(S, G)}{TP(S, G) + FN(S, G)} \quad (3)$$

Algorithm 1 Video superframe clustering algorithm

- 1: **Inputs:**
Number of clusters: K
 - 2: **Initialize:**
 $m = 0.1 * K$ ▷ compactness
Cluster centers $C_k = [x_1 \dots x_f]^T$ at regular step S
 - 3: Perturb cluster centers in a neighborhood, to the lowest gradient position
 - 4: **repeat**
 - 5: $S_t \leftarrow S_{t-1}$
 - 6: **for** each cluster center C_k **do**
 - 7: Assign the best matching frames from a $2S$ neighbourhood around the cluster center according the distance measure (Eq. 1).
 - 8: **end for**
 - 9: Compute new cluster centers and error E {L1 distance between previous centers and recomputer centers}
 - 10: **until** $E \leq$ threshold
 - 11: Postprocessing to remove very short clusters
-

where $TP(S, G)$ is the number of boundary frames in G for which there is a boundary frame in S in range r and $FN(S, G)$ is the number of boundary pixels in G for which there is no boundary frame in S in range r . In simple words, Boundary Recall, Rec , is the fraction of boundary frames in G which are correctly detected in S . In this work, the range r , as a tolerance parameter, is set to 0.008 times the video length in frames based on the experiments.

Under-segmentation error is another error metric which measures the leakage from superframes with respect to the ground truth segmentation. The lower under-segmentation error, the better match between superframes and the ground truth segments. We define under-segmentation error, $UE(S, G)$, as follows:

$$UE(S, G) = \frac{1}{N} \sum_{i=1}^L \left(\sum_{j|s_j \cap g_i > \beta} \min\{|s_j \cap g_i|, |s_j - g_i|\} \right) \quad (4)$$

Where N is the number of frames in the video, L is the number of ground truth superframes, $|\cdot|$ indicates the length of a segment in frames and $S_j - G_i = \{x \in S_j | x \notin G_i\}$. By doing some experiments, we found the best number of β as an overlap threshold, is 0.25 of each superframe s_j .

III. EXPERIMENTS

This proposed algorithm requires two initialization: K (the number of desired clusters) and ‘compactness’. In our work, we initialize compactness to $0.1 * K$ to make the evaluation of results easier.

Experiments are carried out using the ‘MAD’¹ database and the ‘UMN’ databases. The MAD² database [20] is recorded using a Microsoft Kinect sensor in an indoor environment with

¹Multimodal Action Database

²Videos available from www.humansensing.cs.cmu.edu/mad/

a total of 40 video sequences of 20 subjects. Each subject performs 35 sequential actions twice and each video is about 4000–7000 frames. We also labelled superframes manually for each video. That shows in which frames there is a big change of motion for the frame, which illustrates the video segments with similar motions, superframes.

UMN³ is an unusual crowd activity dataset [21], a staged dataset that depicts sparsely populated areas. Normal crowd activity is observed until a specified point in time where behavior rapidly evolves into an escape scenario where each individual runs out of camera view to simulate panic. The dataset comprises 11 separate video samples that start by depicting normal behavior before changing to abnormal. The panic scenario is filmed in three different locations, one indoors and two outdoors. All footage is recorded at a frame rate of 30 frames per second at a resolution of 640×480 using a static camera.

Each frame is represented by 20 features, 11 values for HOM features, 8 values for HOD features, and one value for the frame location in the video. HOF features provide us with a normalized histogram at each frame of the video.

A dense optical flow [1] is used in this work to detect motion for each video frame. We have employed FlowNet-2 which has an improved quality and speed compared to other optical flow algorithms. With the MAD and UMN databases it took only about $50ms$ and $40ms$ respectively per frame to extract the flow using FlowNet-2. Figure 2 illustrates the results of FlowNet-2 on both databases.

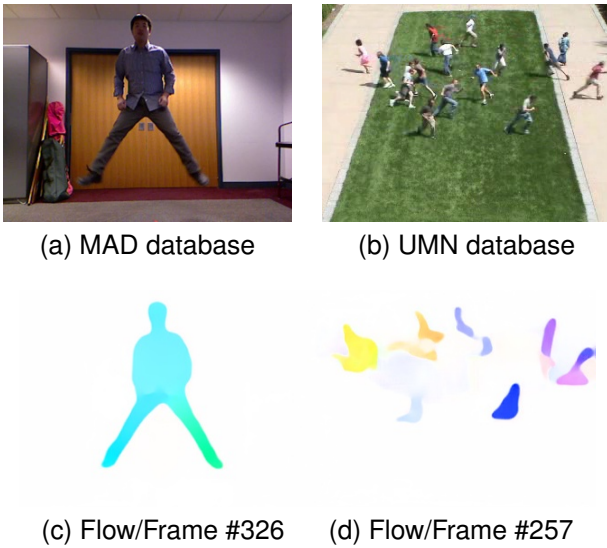


Fig. 2. Examples of optical flow results using FlowNet-2 on a frame of a video from MAD and UMN databases.

To compare our results against, we have used ‘phase correlation’ between two group of frames to determine relative translative movement between them as proposed in [18]. This method relies on estimating the maximum of the phase correlation, which is defined as the inverse Fourier transform

³Videos available from mha.cs.umn.edu/projects.shtml

of the normalized cross-spectrum between two space-time volumes in the video sequence. We call this phase-correlation technique as *PC-clustering* in this work. We sampled every 2nd frame in each video and chose a sub-section of each frame. We use 240×240 pixels in the middle of each frame to carry out the phase correlation. Therefore, each space-time volume is considered to be $240 \times 240 \times 30$, in which 30 is the number of frames and shows the temporal length of each volume.




Fig. 3. The relation between the threshold and the number of clusters in one of the videos of MAD database in PC-clustering technique [18].

Given two space-time volumes v_a and v_b , we calculate the normalized cross-correlation using Fourier transform and taking the complex conjugate of the second result (shown by $*$) and then the location of the peak using equation 6.

$$\mathbf{F}_a = \mathcal{F}(v_a), \quad \mathbf{F}_b = \mathcal{F}(v_b), \quad \mathbf{R} = \frac{\mathbf{F}_a \circ \mathbf{F}_b^*}{|\mathbf{F}_a \circ \mathbf{F}_b^*|} \quad (5)$$

where \circ is the element-wise product.

$$r = \mathcal{F}^{-1}\{\mathbf{R}\}, \quad \mathbf{corr} = \underset{v}{\operatorname{argmax}} \{r\} \quad (6)$$

where \mathbf{corr} is the correlation between two space-time volumes v_a and v_b . We calculate this correlation every 2nd frame of each video and then we conclude that there is a big change in the video when the cross-correlation is less than a threshold (see Figure 3). Figure 4 illustrates how boundary recall and under-segmentation error changes over the number of clusters.

To test the performance of our method on video clustering to superframes, we consider two groups of features: first the averaged value of optical flow over u and v , i.e. the x and y components of the optical flow and second the histogram of magnitude and direction of flow. Figure 5 shows the boundary recall with respect to the input number of desired superframes for one of the videos (*sub01_seq01*) in the MAD dataset. For this video, the number of superframes in the ground truth is 71 which for $K > 70$, HOF features has outstanding improvement on boundary recall over the averaged features.

As discussed before, the output number of clusters is usually less than the input number of desired clusters K . Since in the postprocessing step some of the clusters may get merged with




Fig. 4. Boundary recall and under-segmentation error based on the number of clusters for one of the videos of MAD database for the proposed algorithm.




Fig. 5. Comparison of results when using averaged flow over components of flow and histogram of flow. The former is shown in pink-triangular and the latter is shown in green-circles.

other clusters as their length is too short. Figure. 6 illustrates the relation between the output number of clusters, K and the boundary recall for a video from the MAD database with 71 ground-truth clusters. The ground-truth number of clusters is shown using a red horizontal line in the figure. It is shown that when the output number of clusters for this video is 71 or more, K is more than 115 and the boundary recall is bigger than %91.

The ground-truth and the result boundaries for a video of MAD database are shown in Figure 7. The difference between them is also shown in this figure which is almost zero except for clusters between 50 to 62 between frames 1600 and 2100.

As stated before a combination of under-segmentation error and boundary recall is a good evaluation metric for the algorithm. Under-segmentation error accurately evaluates the quality of superframe segmentation by penalizing superframes overlapping with other superframes. Higher boundary recall also indicates less true boundaries are missed. Figure 8 shows the experimental results for both boundary recall and under-segmentation error which are the average values over all 40

TABLE I
QUANTITATIVE EVALUATION FOR VIDEO CLUSTERING

	MAD database		UMN database	
	UE	Rec	UE	Rec
PC-clustering [18]	0.54	0.64	0.27	0.25
Averaged flow	0.52	0.77	0.19	0.63
Histogram of flow (proposed)	0.31	0.99	0.17	0.71

videos in the MAD database. According to this figure, for K more than about 122, the boundary recall is %100 and the under-segmentation error is less than %50.

A quantitative evaluation of all videos in each dataset is done and the results are illustrated in Table I. These results are calculated by averaging over all the videos in each dataset when the number of clusters is about 115. The experimental results show that our model works quite well on both datasets. Regarding this table, the histogram of flow works as a better feature than just averaging the flow or using phase correlation between volumes in the video. There is a %22 improvement of boundary recall for the MAD dataset and %8 for UMN dataset. The boundary recall is quite low for the baseline on UMN dataset, although the segmentation error is still very low. An interesting point in this table is that, there a higher boundary recall for all methods on the MAD dataset, however lower under-segmentation error on UMN datasets.

IV. CONCLUSION

In this paper, we proposed using Histogram of Optical Flow (HOF) features to cluster video frames. These features are independent of the scale of the moving objects. Using these atomic video segments to speed up later-stage visual processing, has been recently used in some works. The number of desired clusters, K , is the input to our algorithm. This parameter is very important and may cause over/under-segmenting the video.

Our superframe method divides a video into superframes by minimizing a cost function which is distance-based and makes each superframe belong to a single motion pattern and not overlap with others. We used FlowNet-2 to extract motion information per video frame to describe video sequences and quantitatively evaluated our method over two databases, MAD and UMN.

One interesting trajectory for the future work is to estimate the saliency score for each of these superframes. This helps us to rank different episodes of a video by their saliency scores and detect the abnormalities in videos as the most salient segment.

REFERENCES

- [1] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox, "FlowNet 2.0: Evolution of optical flow estimation with deep networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [2] V. Bloom, D. Makris, and V. Argyriou, "Clustered spatio-temporal manifolds for online action recognition," in *International Conference on Pattern Recognition*, 2014, pp. 3963–3968.




Fig. 6. Boundary recall and the number of output clusters for a video from MAD database in blue-circle and orange-square, respectively. The number of superframes in the ground truth for this video is 71 which is shown by a red line.




Fig. 7. Comparison of results when using averaged magnitude/direction and HOF. The number of ground-truth clusters for this video is 71 clusters and the boundary recall is 0.91. The averaged flow results are shown in pink-triangular and the HOF is shown in green-circles.




Fig. 8. Quantitative evaluation of our algorithm by averaging over 40 videos in MAD dataset.

[3] F. Caba Heilbron, V. Escorcia, B. Ghanem, and J. Carlos Niebles, "Activitynet: A large-scale video benchmark for human activity understanding," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 961–970.

[4] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Susstrunk,

"Slic superpixels compared to state-of-the-art superpixel methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 34, no. 11, pp. 2274–2282, 2012.

[5] M. Grundmann, V. Kwatra, M. Han, and I. Essa, "Efficient hierarchical graph-based video segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[6] W.-S. Chu, Y. Song, and A. Jaimes, "Video co-summarization: Video summarization by visual co-occurrence," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 3584–3592.

[7] Q. Tu, A. Men, Z. Jiang, F. Ye, and J. Xu, "Video saliency detection incorporating temporal information in compressed domain," *Image Communication*, vol. 38, no. C, pp. 32–44, 2015.

[8] W. G. Aref, A. C. Catlin, J. Fan, A. K. Elmagarmid, M. A. Hammad, I. F. Ilyas, M. S. Marzouk, and X. Zhu, "A video database management system for advancing video database research," in *International Workshop on Management Information Systems*, 2002, pp. 8–17.

[9] H. A., "Designing video data management systems," Ph.D. dissertation, The University of Michigan, Ann Arbor, MI, USA, 1995.

[10] X. Ren and J. Malik, "Learning a classification model for segmentation," in *IEEE International Conference on Computer Vision (ICCV)*, 2003, pp. 10–17.

[11] M. Reso, J. Jachalsky, B. Rosenhahn, and J. Ostermann, "Temporally consistent superpixels," *IEEE International Conference on Computer Vision (ICCV)*, pp. 385–392, 2013.

[12] J. Chang, D. Wei, and J. W. Fisher III, "A video representation using temporal superpixels," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 2051–2058.

[13] C. Xu and J. J. Corso, "Evaluation of super-voxel methods for early video processing," *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1202–1209, 2012.

[14] S.-H. Lee, W.-D. Jang, and C.-S. Kim, "Temporal superpixels based on proximity-weighted patch matching," in *IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 3610–3618.

[15] N. L. Xiaoquan Yi, "Fast pixel-based video scene change detection," in *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2005.

[16] M. Gygli, "Ridiculously fast shot boundary detection with fully convolutional neural networks," *Computing Research Repository (CoRR)*, vol. abs/1705.08214, 2017.

[17] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *International Journal of Computer Vision (IJCV)*, vol. 59, no. 2, pp. 167–181, 2004.

[18] I. Kotsia and V. Argyriou, "Action spotting exploiting the frequency domain," in *Workshop on CVPR*, 2011, pp. 43–48.

[19] V. Argyriou and T. Vlachos, "Quad-tree motion estimation in the frequency domain using gradient correlation," *IEEE Transactions on Multimedia*, vol. 9, no. 6, pp. 1147–1154, 2007.

[20] D. Huang, S. Yao, Y. Wang, and F. De La Torre, "Sequential max-margin event detectors," in *European Conference on Computer Vision (ECCV)*, 2014, pp. 410–424.

[21] O. a. S. M. Mehran, R., "Abnormal crowd behaviour detection using social force model," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 935–942.