

Detecting Violence in Video using Subclasses

Xirong Li^{1,2,3} Yujia Huo² Qin Jin^{1,2*} Jieping Xu^{1,2}

¹Key Lab of Data Engineering and Knowledge Engineering, Renmin University of China

²Multimedia Computing Lab, School of Information, Renmin University of China

³Shanghai Key Laboratory of Intelligent Information Processing

{xirong,qjin,xjieping}@ruc.edu.cn

ABSTRACT

This paper attacks the challenging problem of violence detection in videos. Different from existing works focusing on combining multi-modal features, we go one step further by adding and exploiting subclasses visually related to violence. We enrich the MediaEval 2015 violence dataset by *manually* labeling violence videos with respect to the subclasses. Such fine-grained annotations not only help understand what have impeded previous efforts on learning to fuse the multi-modal features, but also enhance the generalization ability of the learned fusion to novel test data. The new subclass based solution, with AP of 0.303 and P100 of 0.55 on the MediaEval 2015 test set, outperforms the state-of-the-art. Notice that our solution does not require fine-grained annotations on the test set, so it can be directly applied on novel and fully unlabeled videos. Interestingly, our study shows that motion related features (MBH, HOG and HOF), though being essential part in previous systems, are seemingly dispensable. Data is available at <http://lixirong.net/datasets/mm2016vsd>.

CCS Concepts

•Computing methodologies → Scene understanding;

Keywords

Video violence detection, subclass annotation, fusion

1. INTRODUCTION

Detecting violence in videos is important for protecting children against offensive content. Viewing violence as a specific concept, one might treat violence detection as a showcase of video concept detection, an extensively studied topic in the literature [12, 22]. The MediaEval affective task defines violence as *one would not let an 8 years old child see because of their physical violence* [15]. Scenes of violence lack consistent visual patterns, see Fig. 1. Moreover,

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '16, October 15 - 19, 2016, Amsterdam, Netherlands

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-3603-1/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2964284.2967289>



Figure 1: Visual examples of ten violence subclasses. We enrich the MediaEval 2015 violence data [15] by providing video-level annotations with respect to these subclasses.

the perception of violence varies over subjects. Another evidence showing the different landscape of violence detection is that state-of-the-art video concept detection uses deep image features alone [16]. By contrast, a violence detection system often involves a number of multi-modal features related to image, audio and motion [2, 8, 23]. We also use multi-modal features, but go one step further by adding and exploiting subclasses visually related to violence, for better performance and deeper insight.

Progress on video violence detection. Thanks to the MediaEval benchmark effort [3, 4, 15], data for violence detection research has been substantially expanded in terms of types of violence and the amount of videos, compared to earlier work on a specific type of violence, e.g., fight in hockey games [1]. Consequently, different from previous works that mainly utilize motion related features such as motion SIFT and histogram of optical flow [1, 20], current solutions embrace multi-modal features to describe video content in terms of visual appearance, audio channel and motion over the temporal dimension. A common package, as reported in [2, 6, 8, 23], consists of image features extracted by pre-trained deep convolutional neural networks [14, 17], MFCC based audio features [5], and local descriptors derived from improved dense trajectories [19].

As for combining the many features, although late fusion is consistently found to improve over models derived from single features [6, 7, 23], the success so far has been limited to average weighting, i.e., all the underlying models are treated equally. As the effectiveness of the features varies [8], average fusion is unlikely to be the optimal so-

lution. Nonetheless, efforts on learn-to-fuse are found to be inferior to the average-fusion baseline [6]. While current works simply attribute this to divergence between training and test data [6, 7], we believe a more in-depth analysis is required to understand what impedes the use of learn-to-fuse for video violence detection.

Contributions of this work. In order to better understand the difficulty in generalizing models learned on a given dataset to a novel test set, we first enrich the MediaEval 2015 violence set, an up-to-date benchmark for video violence detection, by providing video-level manual annotations with respect to ten subclasses of violence. We are inspired by the work from Tan and Ngo [18], which makes an interesting attempt to express violence by a number of middle-level concepts. Nonetheless, as they target at learning these concepts from YouTube videos to derive a semantic feature, the two works are driven by completely different motivations. As we will show in Section 2, the newly added subclass annotations help interpret cross-dataset divergence and consequently the difficulty of generalizing the learned fusion to previously unseen test data. Moreover, we show in Section 3 how this difficulty can be resolved to some extent using subclass based violence detection, and learned fusion of multiple features and multiple subclasses outperforms several state-of-the-art alternatives.

2. SUBCLASS ANNOTATION

Video sets. We use the official release of the MediaEval 2015 affective task [15]. Containing 10,900 short video clips extracted from 199 professional and amateur movies, this benchmark dataset is much larger and more diverse than those narrow-domain videos used in previous works [1, 20]. The MediaEval dataset consists of two disjoint subsets: a development set of 6,144 clips and a test set of 4,756 clips. Notice that the videos were not originally selected for violence, so they have low occurrence of violence, with 4.4% in the development set and 4.8% in the test set.

Manual labeling procedure. We aim to provide fine-grained annotations of violence videos by manually assigning visual concepts that are highly relevant to violence. Such a concept is supposed to have relatively consistent visual appearance and thus more easy to be detected. We take the 52 middle-level concepts from [18] as our initial concept vocabulary. Given the subjective definition of video violence, the vocabulary may expand with novel concepts as the manual labeling process goes on. We watched all violence videos in the development set, eventually yielding a vocabulary of 95 concepts. For a specific concept, there shall be a reasonable amount of positive instances, otherwise modeling this concept is very likely to be futile. We empirically set the threshold to be 20, i.e., a concept needs to have over 20 occurrences in the development set. This results in ten subclasses of violence, as exemplified in Fig. 1. Four of them, namely ‘bind’, ‘aim’, ‘weapon’ and ‘death’, are not covered by the initial 52 concepts. A video can be labeled with multiple subclasses. Some subclasses tend to co-occur with some others in the same video clips, e.g., ‘rope’ to ‘bind’ and ‘blood’ to ‘death’. However, ‘death’ and ‘fight’ seldom co-occur, as a fight hardly ends with death in this dataset.

In order to reveal potential divergence between the development set and the test set, we repeat the above procedure to annotate all violence videos in the test set. Fig. 2 shows the occurrence rate of each subclass, which clearly varies

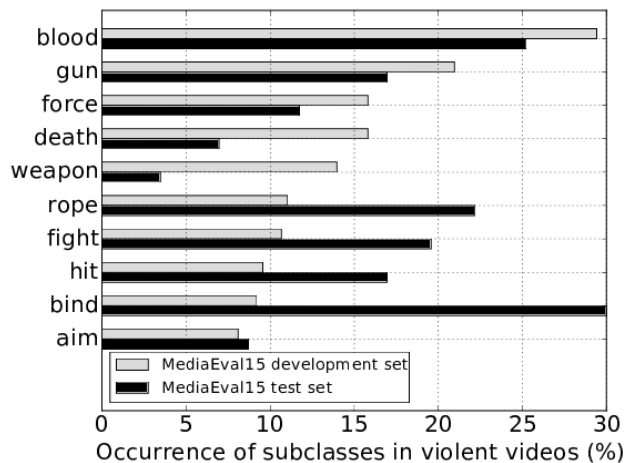


Figure 2: Occurrence of each subclass in violence videos, varying across datasets. Such divergence adds difficulty of generalizing models trained on the development set to the previously unseen test set.

across the two datasets. Such divergence suggests the challenge of applying a fusion model trained on the development set to the (unseen) test set.

Note that the fine-grained annotations of the test set are meant for cross-dataset analytics only. They are not used for violence detection. So, similar to existing works [2, 4, 23], our subclass based violence detection solution is directly applicable to novel video data that are completely unlabeled.

3. USING SUBCLASSES FOR VIOLENCE DETECTION

To verify the viability of subclass based violence detection, we build two systems, with and without subclasses, respectively.

3.1 Setup

Following the current practice [2, 7, 23], we employ multimodal video representations including image features extracted by deep convolutional neural networks [14, 17], audio features [5], and motion features [19]. They are described as follows, and summarized in Table 1.

Image CNN features. Given a video, we uniformly extract its frames with an interval of 0.5 second, resulting in 131,441 and 101,587 frames in the development and test sets, respectively. For each frame we extract deep features using three pretrained CNN models, i.e., VGGNet [14], GoogleNet [17], and GoogLeNet4k [11], a very recent variant of GoogleNet trained on a bottom-up reorganization of the ImageNet hierarchy. By taking the last fully connected layer of VGGNet and the pool5 layer of the two GoogleNets, we obtain three frame-level features

A video-level representation substantially reduces the amount of training/test examples and is commonly used in event detection [21, 22]. We obtain three video-level features by average pooling on the corresponding frame-level features.

Audio features. We use the Mel-frequency Cepstral Coefficients (MFCCs) [5], computed over a sliding short-time window of 25ms with a 10ms shift. An audio segment is then

Table 1: Fourteen features used in our experiments, describing video content in varied aspects.

Modality	Feature	Notation	Dim.
Image	frame-level vggnet	vnet _f	4,096
	video-level vggnet	vnet _v	4,096
	frame-level googlenet	gnet _f	1,024
	video-level googlenet	gnet _v	1,024
	frame-level googlenet4k	g4k _f	1,024
	video-level googlenet4k	g4k _v	1,024
Audio	mfcc + bow	mfcc _b	4,096
	mfcc + fisher vector	mfcc _{f_v}	19,968
Motion	mbh + bow	mbh _b	4,000
	mbh + fisher vector	mbh _{f_v}	98,304
	hog + bow	hog _b	4,000
	hog + fisher vector	hog _{f_v}	49,152
	hof + bow	hof _b	4,000
	hof + fisher vector	hof _{f_v}	55,296

represented by a set of MFCC features. In order to transform these features into a fixed-dimension vector, we adopt two encoding schemes, i.e., bag-of-audio-words by hard assignment and Fisher Vector encoding [13].

Motion features. The Improved Dense Trajectory (IDT) [19] is used to capture actions in a video. We use three trajectory based descriptors, namely Motion Boundary Histogram (MBH), Histogram of Oriented Gradient (HOG), and Histogram of Optical Flow (HOF). In a similar manner to the audio features, bag-of-words and Fisher Vector encodings are used to quantize each of the descriptors, resulting in six motion related features.

Classification Models. For each of the ten subclasses as well as the holistic violence class, we train binary linear SVMs classifiers. To effectively learn from the large amounts of imbalanced examples, we leverage the Negative Bootstrap algorithm [9], which iteratively and adaptively finds negative training examples that better improve classification than negatives selected at random. For each frame-level classifier, we obtain its prediction on a given video by first scoring every frame, smoothing the scores along the temporal dimension, and then taking the maximal response.

The development set is randomly divided into two disjoint subsets, 70% for training the SVM models and 30% as a validation set to optimize hyper-parameters.

Strategies for multi-feature multi-subclass fusion. We investigate late fusion with two weighting strategies. One is average weighting, i.e., assigning equal weights to the base classifiers. Despite its simplicity, the average strategy is found to be effective for multi-feature fusion for violence detection [8]. The second strategy is learning to fuse. In particular, we optimize the weights on the validation set using coordinate ascent, which optimizes a chosen weight per iteration. A recent study in the context of image classification [10] shows that compared to more complicated fusion algorithms such as multiple kernel learning, coordi-

Table 2: Performance of the individual features, sorted in descending order by AP scores of the *learn* run on the test set. For each feature, top performers on a given dataset are shown in bold font.

Feature	MediaEval15 validation set			MediaEval15 test set		
	<i>w/o subclass</i>	<i>avg</i>	<i>learn</i>	<i>w/o subclass</i>	<i>avg</i>	<i>learn</i>
g4k _f	0.321	0.260	0.365	0.190	0.179	0.216
g4k _v	0.284	0.237	0.292	0.213	0.195	0.199
vnet _v	0.295	0.263	0.316	0.120	0.137	0.156
gnet _f	0.245	0.210	0.316	0.097	0.164	0.156
hog _b	0.164	0.132	0.174	0.160	0.175	0.153
vnet _f	0.347	0.226	0.339	0.118	0.147	0.145
mfcc _{f_v}	0.313	0.270	0.316	0.100	0.138	0.136
gnet _v	0.296	0.258	0.312	0.152	0.142	0.130
mbh _b	0.119	0.109	0.197	0.122	0.129	0.092
hof _b	0.127	0.130	0.170	0.113	0.109	0.092
mfcc _b	0.320	0.234	0.295	0.090	0.075	0.077
mbh _{f_v}	0.199	0.075	0.159	0.099	0.064	0.067
hof _{f_v}	0.134	0.083	0.096	0.102	0.050	0.054
hog _{f_v}	0.225	0.070	0.101	0.078	0.047	0.049

nate ascent is more effective by directly optimizing (non-differential) performance metrics.

Putting all the features and subclasses together, we need to learn weights for up to 140 classifiers. By contrast, for the system without subclasses, we only need to optimize weights for 14 classifiers at maximum. Hence, comparing the two systems shall allow us to conclude whether the introduction of subclasses makes sense.

Evaluation criteria. Following the protocol [15], we report Average Precision (AP), Precision at the top 10 ranked videos (P10), and P100.

3.2 Experiments

Experiment 1. The impact of subclasses in a single-feature setting. Table 2 shows the performance of the individual features, where *w/o class* indicates runs without considering subclasses, *avg* means average fusion of the ten subclasses, and *learn* is fusion with weights tuned on the validation set. In the *w/o class* setting, the best motion feature hog_b, with AP of 0.160, outperforms the best standard CNN feature, i.e., gnet_v. This is consistent with a recent result reported in [8]. Nonetheless, g4k_v with AP of 0.213 is the top performer.

Half of the features improve by combining subclass classifiers, with the best performance obtained by gnet_f + *learn*. Interestingly, the majority of the features for which fusion does not work are motion related. By looking in combinations of one subclass and one feature for violence detection, we find that the motion features are in general less effective for modeling subclasses when compared with image CNN features.

Experiment 2. The impact of subclasses in a multi-feature setting. We now move to the setting where multiple features are combined together with subclasses. Given that there are over 16k combinations, we take a more practical approach by starting with using all the features. Specific

Table 3: Performance of feature-subclass fusion, sorted in descending order according to AP scores of the *learn* run on the test set. For each feature-fusion setting, top performers on a given dataset are shown in bold font.

Varied settings of feature fusion														MediaEval15 val set				MediaEval15 test set			
Image CNN						Audio MFCC		Motion IDT						w/o subclass		with subclass		w/o subclass		with subclass	
vnet _f	vnet _v	gnet _f	gnet _v	g4k _f	g4k _v	mfcc _b	mfcc _{f_v}	mbh _b	mbh _{f_v}	hog _b	hog _{f_v}	hof _b	hof _{f_v}	<i>avg</i>	<i>learn</i>	<i>avg</i>	<i>learn</i>	<i>avg</i>	<i>learn</i>	<i>avg</i>	<i>learn</i>
✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	0.484	0.511	0.386	0.510	0.274	0.251	0.288	0.303
✗	✗	✗	✗	✓	✓	✓	✓	✓	✗	✓	✗	✓	✗	0.376	0.492	0.293	0.493	0.269	0.275	0.277	0.296
✗	✗	✗	✗	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	0.455	0.473	0.397	0.491	0.248	0.236	0.283	0.290
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	0.442	0.531	0.279	0.522	0.248	0.270	0.226	0.279
✓	✓	✓	✓	✓	✓	✓	✗	✓	✗	✓	✗	✓	✗	0.440	0.527	0.282	0.511	0.300	0.260	0.259	0.275
✓	✓	✓	✓	✓	✓	✗	✗	✓	✗	✓	✗	✓	✗	0.394	0.482	0.254	0.435	0.280	0.225	0.249	0.263
✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	0.468	0.495	0.353	0.512	0.266	0.228	0.250	0.263
✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	0.427	0.449	0.291	0.419	0.244	0.233	0.239	0.255
✓	✓	✓	✓	✓	✓	✗	✗	✓	✓	✓	✓	✓	✓	0.373	0.479	0.226	0.447	0.229	0.249	0.202	0.253
✓	✓	✓	✓	✓	✓	✗	✗	✗	✓	✗	✓	✗	✓	0.434	0.468	0.250	0.444	0.238	0.253	0.171	0.231
✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	0.308	0.424	0.194	0.446	0.164	0.141	0.141	0.141
✗	✗	✗	✗	✗	✗	✓	✓	✗	✗	✗	✗	✗	✗	0.347	0.348	0.273	0.353	0.107	0.106	0.111	0.117
✗	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓	✓	0.206	0.279	0.146	0.296	0.143	0.111	0.116	0.114

features are progressively removed with preference given to slower and longer features. E.g., fisher vectors are removed in advance to their bag-of-words counterparts. Moreover, in order to study complementariness of distinct modalities, features of an unchosen modality are left out. Eventually we obtain 13 feature combinations with their performance reported in Table 3.

Comparing with the single feature setting, multi-feature fusion improves the performance. In the *w/o subclass* setting, however, fusion weights learned on the validation set often yields lower AP than average weighting. As shown in Table 3, for only 4 out of 13 times *learn* is better than *avg*. By contrast, the use of subclasses makes *learn* beats *avg* 11 out of 13 times. This result shows the importance of subclasses for learning to fuse.

We observe that some subclasses are more easier to detect, e.g., ‘blood’, ‘gun’ and ‘death’, than some others, e.g., ‘aim’, ‘rope’, and ‘bind’. Due to page limit, more detailed analysis concerning the subclass detection will be presented in the future extension of this work.

Concerning the complementariness of distinct modalities, all the modalities are required for deriving a holistic violence detector, as in the *w/o subclass* scenario. However, once subclasses are taken into account, the joint use of image and audio modalities (with motion excluded) performs the best.

Finally, we make a system level comparison between the proposed subclass based system and three state-of-art systems as demonstrated in the MediaEval 2015 evaluation. The result shown in Table 4 further confirms the effectiveness of subclass based violence detection.

4. SUMMARY AND CONCLUSIONS

We contribute to video violence detection as follows. First, we enrich the MediaEval 2015 violence dataset by providing video-level annotations with respect to ten subclasses of violence. The clear divergence between subclass distributions of the MediaEval dev and test sets explains why existing learn-to-fuse attempts did not work. Second, by a

Table 4: Comparing with front-runners in the MediaEval 2015 violence detection task.

System	AP	P10	P100
Fudan-Huawei [2]	0.296	1.0	0.46
MIC-TJU [23]	0.285	0.8	0.49
NII-UIT [7]	0.268	0.6	0.44
<i>this work, w/o subclass</i>	0.275	0.9	0.42
<i>this work, with subclass</i>	0.303	1.0	0.55

systematic comparison between violence detection with and without subclasses, we re-justify the effectiveness of learn-to-fuse for violence detection. The new solution beats the state-of-the-art on the MediaEval 2015 test set. Moreover, decomposing violence into subclasses results in novel findings concerning complementariness of distinct modalities. The conclusion of [8] that visual and motion features are better than audio features holds when modeling violence in a holistic manner. For fusion of multiple subclasses and multi-modal features, motion related features appear to be ignorable.

Acknowledgments. The authors are grateful to the MediaEval 2015 Affective Impact of Movies Task organizers [15] for making video data available for this research. The work was supported by the Fundamental Research Funds for the Central Universities and the Research Funds of Renmin University of China (No. 14XNLQ01), NSFC (No. 61303184, No. 71531012), SRFDP (No. 20130004120006), BJNSF (No. 4142029), and Shanghai Key Laboratory of Intelligent Information Processing, China (Grant No. I IPL-2013-002).

5. REFERENCES

- [1] E. Bermejo, O. Deniz, G. Bueno, and R. Sukthankar. Violence detection in video using computer vision techniques. In *CAIP*, 2011.

- [2] Q. Dai, R. Zhao, Z. Wu, X. Wang, Z. Gu, W. Wu, and Y. Jiang. Fudan-Huawei at MediaEval 2015: Detecting Violent Scenes and Affective Impact in Movies with Deep Learning. In *MediaEval Workshop*, 2015.
- [3] C. Demarty, C. Penet, M. Schedl, B. Ionescu, V. L. Quang, and Y. Jiang. The MediaEval 2013 Affect Task: Violent Scenes Detection. In *MediaEval Workshop*, 2013.
- [4] C. Demarty, C. Penet, M. Soleymani, and G. Gravier. VSD, a public dataset for the detection of violent scenes in movies: design, annotation, analysis and evaluation. *Multimedia Tools Appl.*, 74(17):7379–7404, 2015.
- [5] T. Ganchev, N. Fakotakis, and G. Kokkinakis. Comparative evaluation of various MFCC implementations on the speaker verification task. In *SPECOM*, 2005.
- [6] Q. Jin, X. Li, H. Cao, Y. Huo, S. Liao, G. Yang, and J. Xu. RUCMM at MediaEval 2015 affective impact of movies task: Fusion of audio and visual cues. In *MediaEval Workshop*, 2015.
- [7] V. Lam, S. P. Le, D. Le, S. Satoh, and D. A. Duong. NII-UIT at MediaEval 2015: Affective Impact of Movies Task. In *MediaEval Workshop*, 2015.
- [8] V. Lam, S. Phan, D.-D. Le, D. A. Duong, and S. Satoh. Evaluation of multiple features for violent scenes detection. *Multimedia Tools Appl.*, 2016. in press.
- [9] X. Li, C. G. Snoek, M. Worring, D. Koelma, and A. W. Smeulders. Bootstrapping visual categorization with relevant negatives. *IEEE Trans. Multimedia*, 15(4):933–945, 2013.
- [10] S. Liao, X. Li, H. T. Shen, Y. Yang, and X. Du. Tag features for geo-aware image classification. *IEEE Trans. Multimedia*, 17(7):1058–1067, 2015.
- [11] P. Mettes, D. C. Koelma, and C. G. M. Snoek. The ImageNet Shuffle: Reorganized Pre-training for Video Event Detection. In *ICMR*, 2016.
- [12] M. Naphade, J. R. Smith, J. Tesic, S.-F. Chang, W. Hsu, L. Kennedy, A. Hauptmann, and J. Curtis. Large-scale concept ontology for multimedia. *IEEE MultiMedia*, 13(3):86–91, 2006.
- [13] J. Sánchez, F. Perronnin, T. Mensink, and J. J. Verbeek. Image classification with the fisher vector: Theory and practice. *International Journal of Computer Vision*, 105(3):222–245, 2013.
- [14] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*, 2015.
- [15] M. Sjöberg, Y. Baveye, H. Wang, V. L. Quang, B. Ionescu, E. Dellandrea, M. Schedl, C.-H. Demarty, and L. Chen. The MediaEval 2015 affective impact of movies task. In *MediaEval Workshop*, 2015.
- [16] C. Snoek, S. Cappallo, D. Fontijne, D. Julian, D. Koelma, P. Mettes, K. van de Sande, A. Sarah, H. Stokman, and R. Towal. Qualcomm Research and University of Amsterdam at TRECVID 2015: Recognizing Concepts, Objects, and Events in Video. In *NIST TRECVID Workshop*, 2015.
- [17] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [18] C. C. Tan and C.-W. Ngo. The Vireo team at MediaEval 2013: Violent scenes detection by mid-level concepts learnt from Youtube. In *MediaEval Workshop*, 2013.
- [19] H. Wang and C. Schmid. Action Recognition with Improved Trajectories. In *ICCV*, 2013.
- [20] L. Xu, C. Gong, J. Yang, Q. Wu, and L. Yao. Violent video detection based on MoSIFT feature and sparse coding. In *ICASSP*, 2014.
- [21] Z. Xu, Y. Yang, and A. Hauptmann. A discriminative CNN video representation for event detection. In *CVPR*, 2015.
- [22] G. Ye, Y. Li, H. Xu, D. Liu, and S.-F. Chang. EventNet: A large scale structured concept library for complex event detection in video. In *ACMMM*, 2015.
- [23] Y. Yi, H. Wang, B. Zhang, and J. Yu. MIC-TJU in MediaEval 2015 Affective Impact of Movies Task. In *MediaEval Workshop*, 2015.