

HIGHLIGHTS EXTRACTION FROM SPORTS VIDEO BASED ON AN AUDIO-VISUAL MARKER DETECTION FRAMEWORK

Ziyou Xiong[†], Regunathan Radhakrishnan[‡], Ajay Divakaran[‡] and Thomas S. Huang[†]

[†]Department of Electrical and Computer Engineering,

University of Illinois at Urbana-Champaign, Urbana, IL, USA

[‡]Mitsubishi Electric Research Laboratories, Cambridge, MA, USA

E-mail: {zxiong, huang}@ifp.uiuc.edu, {regu, ajayd}@merl.com

ABSTRACT

We propose to use a visual object (e.g., the baseball catcher) detection algorithm to find local, semantic objects in video frames in addition to an audio classification algorithm to find semantic audio objects in the audio track for sports highlights extraction. The highlight candidates are then further grouped into finer-resolution highlight segments, using color or motion information. During the grouping phase, many of the false alarms can be correctly identified and eliminated. Our experimental results with baseball, soccer and golf video are promising.

Keywords: Sports Highlights, Audio-Visual Markers, Finer-Resolution Highlights

1. INTRODUCTION AND RELATED WORK

Sports highlights extraction is one of the most important applications of video analysis. Approaches based on audio classification[1], video feature extraction and highlights modelling[2] have been reported. Due to the space limitation, please refer to the introduction section of [1][2][3] for a detailed literature survey.

Since the content of sports video is intrinsically multimodal, many researchers have also proposed different information fusion schemes to combine different modality information. We categorize some reported fusion methods for sports video analysis as follows:

Simultaneous or Sequential fusion. Hanjalic [4] modelled audience excitement using a function of the following three factors from different modalities: motion activity, density of cuts, and audio energy. He derived an “excitement” function in terms of these three parameters in a simultaneous fashion. Duan et al. [5] used excited speech and goalpost detection together to detect soccer highlights. On the other hand, Chang et al. [6] primarily used audio analysis as an alternative tool for sports parsing. Their goal was to detect football touchdowns. Vision-based line-mark and goal-posts detection were then used to verify the results

obtained from audio analysis.

Statistical or knowledge-based fusion. For the statistical fusion, in [7], Huang et al. have compared 4 different Hidden Markov Models (HMMs) based methods for the task of differentiating advertisement, basketball, football, news, and weather forecast video. For knowledge-based fusion, Rui et al. [1] use a weighted sum of likelihood to fuse the excited speech likelihood and ball-bat impact likelihood. The weight factors are derived from the a priori knowledge regarding which factor receives larger weights.

2. OUR MOTIVATION

In this paper we explore the possibility of building a unified framework to extract highlights from three kinds of sports, i.e., baseball, golf and soccer. Computational power constraints on devices such as TiVo, WebTV, and personal digital video recorders rule out having a completely distinct highlights extraction algorithm for each sport, and thus motivate us to look for general features that would work across different sports. Different from Hanjalic’s approach [4] that relies on low level features, we want to look for semantically higher level features from audio visual content. We summarize our framework in Fig. 1.

3. PROPOSED FRAMEWORK

3.1. Visual Markers Detection

Instead of searching for motion activity pattern or cut density pattern, we target key visual objects that are highly correlated with the interesting segments of the games. For example, for baseball games, we want to detect the pattern in which catcher squats waiting for the pitcher pitching the ball; for golf games, we want to detect the players bending to hit the golf ball; for soccer, we want to detect the appearance of the goal post. Correct detection of these key visual objects can eliminate the majority of the video content that is not in the vicinity of the interesting segments. For the

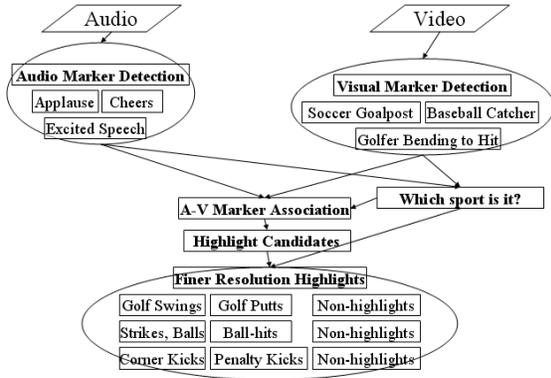


Fig. 1. Our proposed framework: an overview.

goal of one general framework for all three sports, we use the following processing strategy: for the unknown sports content, we detect whether there are baseball catchers, or golfers bending to hit the ball, or soccer goal posts. The detection results can enable us to decide which sport (baseball, golf or soccer) it is.

3.2. Audio-Visual Markers Association

Ideally each visual marker can be associated with one and only one audio marker and vice versa. But, an audio marker without any visual marker in the vicinity may indicate a false alarm and vice versa. Thus a post-processing step, which we call audio-visual marker association, will have the advantage of reducing false alarms.

3.3. Finer-resolution Highlights Recognition

Highlight candidates delimited by the audio markers and visual markers are quite diverse. For example, golf swings and putts share the same audio markers (audience applause and cheering) and visual markers (golfers bending to hit the ball). Both of these two kinds of golf highlight events can be found by the aforementioned audio-visual markers detection based method. Note that using low-level features at the very beginning before audio-visual marker association will not be very useful because the concept of “sports highlights” is at such a high semantic level that it is difficult for the low level features to map to. However, after picking highlight candidates using audio-visual markers we are able to use low-level features for post-processing with success, as can be seen in Table 2.

4. DETAILED METHODS AND RESULTS

We use Gaussian mixture models to classify the sports audio. Due to space limit, we omit the details. Please see [3].

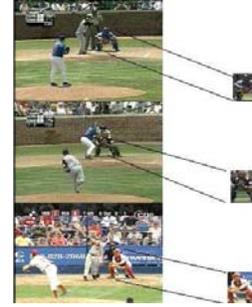


Fig. 2. Some examples of the typical view of the squatting baseball catcher.



Fig. 3. The first few weak classifiers learned by the AdaBoost algorithm for the catcher model.

4.1. Baseball Catcher Detection

We make use of the following observation of the typical baseball video content: at the beginning of the baseball pitch, the camera angle is in such a way that the batter and the catcher almost always face the viewers and we almost always see the frontal, upright view of the catcher squatting to catch the ball. Some examples of these images are shown in Fig. 2.

We use an object detection algorithm such as Viola and Jones’s [8] to detect baseball catchers from baseball video. We have collected more than 8000 thumbnail catcher images (of size 24 x 24) such as those in Fig. 2 as positive examples. We have also collected more than 1000 images of size 352 x 240 from various baseball games. None of these images has any catcher in them. Millions of negative examples can be extracted from these images.

Like the first few features for face detection in [8], the initial rectangle features selected by AdaBoost based on the training baseball catcher images and negative examples are also meaningful and easily interpreted. We show the first 3 such features in Fig. 3. The first feature looks at the difference between “the sum of part of the leg region and the chest region” and “the sum of the hand region and the background region below the butt.” The second feature reflects the difference between the chest region and the sum of the two arm regions. The third feature tells the difference between the very bottom portion of the image that has only the play field region (grass, soil, etc.) and the region right above it that has the player’s feet.

The learned catcher model is used to detect catchers

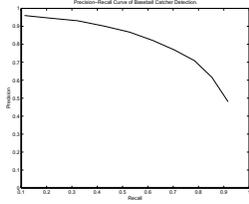


Fig. 4. The precision-recall curve of baseball catcher detection. Y axis: precision; X axis: recall.

from all the video frames in a 3-hour baseball game (Atlanta-Boston, MLB 2002).

We then use the following technique to eliminate some sparse detections: for every frame, we look at a range of frames (from the 14th frames before it to the 14th frames after it), if the number of frames that have had detection declared is above a threshold, then we declare this frame as a frame that has detection. Otherwise, we declare this frame as a frame that does not have any detection. By varying this threshold (a percentage of the total number of frames in the range, in this case, 29), we can compare the number of detections with those in the ground truth set (marked by human viewers). We show the precision-recall curve in Fig. 4 and the detailed results in Table 1.

Table 1. Precision-recall values for catcher detection.

Threshold	Pr	Re	Threshold	Pr	Re
0.1	0.480	0.917	0.2	0.616	0.853
0.3	0.709	0.784	0.4	0.769	0.704
0.5	0.832	0.619	0.6	0.867	0.528
0.7	0.901	0.428	0.8	0.930	0.323
0.9	0.947	0.205	1.0	0.960	0.113

4.2. Soccer Goalpost Detection

We make use of the following two observations from soccer videos. When most of the interesting plays such as goals, corner kicks, penalty kicks take place, the goal post is almost always in the view. Hence detection of the goal post can bring us to the vicinity of these interesting plays with high accuracy. There are mainly two views of the goal post that we need to detect. The cameras are usually positioned in the center of the two sides of the field. The cameramen use pan to go back and forth between two halves of the fields and zoom to focus on special targets. Since the distance between the camera and either of the two goalposts is relatively much larger than the size of the goalpost itself, there is little amount of change of the pose of the goalpost during

the entire process of a game, irrespective of the camera pan or zoom. Some of the example images are shown in Fig. 5.

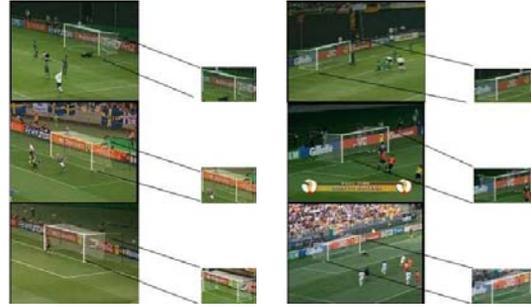


Fig. 5. Some examples of the typical two views of the goalpost.

Detection results of either of the two goalpost are comparable to those of baseball catchers. We omit the details here.

4.3. Golfer Detection

Since there are many different views of the golfer in unconstrained outdoor lighting and camera angle conditions, we have not been able to achieve high detection rate with low false alarm. Fortunately, because golf audio is exceptionally easy to classify compared to baseball and soccer audio, we choose instead to rely on the audio cues in the golf video for highlights extraction. Please refer [3] for details.

Our audio-visual marker tool chest has the flexibility to operate in an audio-only or video-only mode if necessary.

4.4. Audio-Visual Marker Association

If a contiguous sequence of visual markers overlaps with a contiguous sequence of audio markers by a large margin (e.g., the percentage of overlapping is greater than 50%), then we form a “highlight” segment spanning from the beginning of the visual marker sequence to the end of the audio visual marker sequence. Otherwise, we associate a visual marker sequence with the nearest audio marker sequence that follows it if the duration between the two is less than a duration threshold (e.g., the average duration of a set of training “highlights” clips from baseball games).

4.5. Finer-resolution Highlights Classification

The diverse baseball highlight candidates found after the audio markers and visual markers negotiation step are further separated using the techniques described here. For baseball, there are two major categories of highlight candidates, the first being “balls or strikes” in which the batter does not hit the ball, the second being “ball-hits” in which the ball is hit

Table 2. Results before and after the finer-resolution highlights classification.

Video Length	540 min	
Method	A-V Association	Finer-Resolution
# of Candidates	205	173
# of False Alarms	50	18
Highlight Length	29.5 min	25.8 min

to the field or audience. These two categories have different color patterns. In the first category, the camera is fixed at the pitch scene, so the variance of color distribution over time is low. In the second category, in contrast, the camera first shoots at the pitch scene, then it follows the ball to the field or the audience, so the variance of color distribution over time is higher.

We extract the 16-bin color histogram (using the Hue component in the HSV color space) from every video frame of each of the highlight candidate video clip. We use the following algorithm to do the finer-resolution highlights classification:

1. For every color histogram matrix, calculate the “clip-level” mean vector and the “clip-level” standard deviation (STD) vector over its rows.
2. Cluster all the highlight candidate video clips based on their “clip-level” STD vectors into 2 clusters.
3. For each of the clusters,
 - 3.1. Calculate the “cluster-level” mean vector and the “cluster-level” STD vector over the rows of the all the color histogram matrices within the cluster.
 - 3.2. If the value at any color bin of the “clip-level” mean vector is outside the 3σ range of the “cluster-level” mean vector where σ is the STD of the “cluster-level” STD vector at the corresponding color bin, remove it from the highlight candidate list.

When testing on the baseball game, of all the 205 highlight candidates, we have removed 32 using the above algorithm. These 32 clips have been further confirmed to be false alarms by human viewers. The results of the baseball game is shown in Table 2.

5. DISCUSSIONS AND CONCLUSIONS

We have used the visual marker detection (e.g the baseball catcher) algorithm to find local, semantic objects in video frames. In parallel, we have also used audio classification algorithm to find semantic audio objects in the audio track. The highlight candidates are then further grouped into finer-resolution ones, using color or motion information. During the grouping phase, many of the false alarms can also be correctly identified and eliminated.

The framework proposed in this paper is an extension of many of the “audio-visual event analysis” works that are reviewed in [1][2][3] where they can be grouped into a paradigm of “audio-visual features extraction + classifiers”. Our approach adds an important component, i.e., audio-visual marker detection, to this paradigm. Although visual object detection has its own challenges (multi-view object detection is still an active area of research), we have shown it is feasible to concentrate on a single view (in the baseball catcher case) or two views (in the soccer goalpost detection case) for some types of sports.

6. REFERENCES

- [1] Y. Rui, A. Gupta, and A. Acero, “Automatically extracting highlights for TV baseball programs,” in *Proceedings of the Eighth ACM International Conference on Multimedia*, 2000, pp. 105–115.
- [2] L. Xie, S. F. Chang, A. Divakaran, and H. Sun, “Structure analysis of soccer video with hidden Markov models,” in *Proceedings of the International Conference on Acoustic, Speech, and Signal Processing*, vol. 4, May 2002, pp. 4096–4099.
- [3] Z. Xiong, R. Radhakrishnan, and A. Divakaran, “Effective and efficient sports highlights extraction using the minimum description length criterion in selecting GMM structures,” in *Proceedings of the International Conference on Multimedia and Expo*, June 2004.
- [4] A. Hanjalic, “Generic approach to highlight detection in a sport video,” in *Proceedings of the IEEE International Conference on Image Processing*, vol. 1, Sept. 2003, pp. 1–4.
- [5] L.-Y. Duan, M. Xu, T.-S. Chua, Q. Tian, and C.-S. Xu, “A mid-level representation framework for semantic sports video analysis,” in *Proceedings of ACM Conference on Multimedia*, Nov. 2003, pp. 33–44.
- [6] Y.-L. Chang, W. Zeng, I. Kamel, and R. Alonso, “Integrated image and speech analysis for content-based video indexing,” in *Proceedings of the IEEE International Conference Multimedia Computing and Systems*, June 1996, pp. 306–313.
- [7] J. Huang, Z. Liu, Y. Wang, Y. Chen, and E. K. Wong, “Integration of multimodal features for video scene classification based on HMM,” in *Proceedings of IEEE Third Workshop on Multimedia Signal Processing*, Sept. 1999, pp. 53–58.
- [8] P. Viola and M. Jones, “Robust real-time object detection,” *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, May 2004.