

COMPOSING PERSONALIZED VIDEO PLAYOUTS USING SEARCH

Dag Johansen¹, Håvard Johansen¹, Pål Halvorsen^{4,5}, Bjørn Olstad², Cathal Gurrin^{1,3}, Carsten Griwodz^{4,5}

¹University of Tromsø, Norway ²Microsoft ³Dublin City University, Ireland

⁴University of Oslo, Norway ⁵Simula Research Laboratory, Norway

Emails: dag@cs.uit.no, haavardj@cs.uit.no, paalh@ifi.uio.no,
bjornol@microsoft.com, cgurrin@computing.dcu.ie, griff@ifi.uio.no

ABSTRACT

We conjecture that composition of video events from various sources into personalized video playouts will become an important part of next generation streaming systems. Here, video search is a key component since it enables users to retrieve candidate video events based on their interests. One of the main challenges, however, is to analyze the videos in order to correctly identify the various events used to annotate and index the video data. Key problems with current video analysis solutions include that they 1) are complex and therefore require a lot of processing time resulting in large delays; 2) that they can only identify a limited set of events; and 3) that they are still too inaccurate, both giving false positives and failing to find all events. In our DAVVI prototype, we therefore extract metadata for our video search engine by combining existing automatic video analysis tools with currently untapped textual information available in the Internet. This provides an end-user experience where textual query results can be combined dynamically into seamless, highly personalized video playouts using an adaptive torrent-like HTTP streaming solution.

Keywords— Personalized video playout, search, live text commentary metadata

1. INTRODUCTION

YouTube is a streaming service example that empowers end-users. Everybody can produce and make public short video snippets, and everybody can easily determine if, when, and what to download. We consider this extremely popular Internet service a first step away from the traditional multimedia entertainment production and dissemination process with few producers and many consumers. In this context, we are investigating the next step including elements like even stronger user involvement and extreme personalization.

A quickly emerging service is sports streaming, both live and on-demand, where users may retrieve whole events or predefined highlights, and where the video data is presented with all kind of statistics about players, games, and leagues. Examples include American football, basketball, hockey, soccer, and

rugby. As a representative example in this application domain, we have chosen soccer. Several professional content providers deliver popular real-time soccer streaming services, like Schibsted's VG Live and TV2 Sumo in Norway. At best, these systems provide manually generated (static) video summaries of individual games. Thus, they do neither encourage nor enable users to create their own compositions. Also recommendations of a user's own video playouts and video summaries are virtually non-existing in these services.

We conjecture that we will get a closer integration of traditional streaming systems with end-user uploading, sharing, and delivery of personalized content in the near future. As a first demonstrator of such a system, we have built DAVVI [1], a scalable prototype integrating traditional soccer video streaming with video search, recommendation, and personalization. End-users may use the extended functionality to dynamically generate a personalized video (playlist) that is played out as one continuous video whose quality adapts to available resources. Search is a key enabler for such an application, and users can efficiently search and retrieve highlights and combine arbitrary events in a customized manner (like using drag-and-drop) to generate a continuous, personalized video playout on-the-fly from various video sources.

In this paper, we briefly describe our prototype with a novel integration of well-known and new components. In particular, we address the major challenge of providing users with an efficient and accurate search tool that goes far beyond the search functionality available in existing services like YouTube.

2. DAVVI

The DAVVI system [1] provides the functionality required for on-the-fly composition of topic-based, personalized videos that may be played out as one continuous video. As shown in Figure 1, the functional components can be divided into two parts: those that deal with video delivery and those that deal with video search and composition. The services provided by each part are integrated by the client-side software.

DAVVI uses a segmented, torrent-like HTTP streaming approach where similar systems have been shown to be very scalable. Such a solution has also been adopted by, for instance, Move Networks, Microsoft's Smooth Streaming, and Apple's

This research is sponsored by the iAD research based innovation centre, project number 174867 at the Norwegian Research Council.

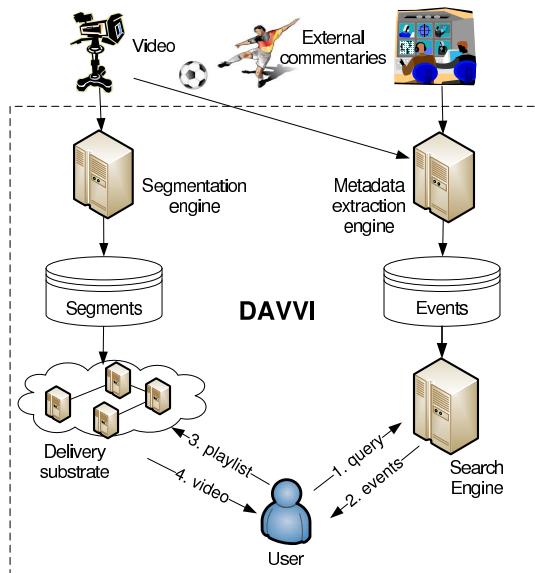


Fig. 1. Architectural overview of DAVVI

HTTP Live streaming. In DAVVI, each video is first processed by the *segmentation engine*, which splits the video into a set of fixed size segments. A video scene usually consists of several consecutive segments. Based on experiences concerning delivery efficiency, coding efficiency, and search result granularity, and with respect to user-accepted quality adaptation frequency [2], our solution uses two-second video segments.

Each segment is encoded as a small self-contained video, which is a different approach than existing systems. This encoding scheme enables composition of multiple segments from different videos into arbitrary playback order. For quality adaptation to available resources, each segment is also encoded in several quality levels (and thus bit rates). To enable such functionality, we currently use closed groups of pictures with H.264 video and MP3 audio, but to 1) reduce start-up and jump latencies, 2) overcome timestamp problems with `libavcodec` which we use for encoding, and 3) avoid the large MPEG2 transport stream container overhead, the media data is packed in a more efficient custom-made container [3].

Each of the segments is then stored in DAVVI's distributed segment store and made available over the Internet through standard web-servers. A tracker holds information about the location of each segment. Based on this information, the client player retrieves video segments using HTTP GET requests. On the client side, resource availability is taken into account when selecting the video quality of the segments to retrieve. Furthermore, the arbitrarily ordered playlist is seamlessly played out by concatenating segments from one or more videos.

Metadata is extracted from external textual data sources by the *metadata extraction engine* and stored within DAVVI as a collection of independent video *events*. Segmentation and annotation techniques based on utilizing such sources have been shown to provide by far the most reliable method for event detection of sports content with almost 100% precision and re-

call when a suitable high-quality source of annotations is available [4]. Furthermore, each event contains a unique video identifier, the event start time, a duration, and a textual event annotation. DAVVI's *search engine* component, using either Apache's Solr or Microsoft's Search Server Express (both of which have successfully been employed), takes the set of identified events as input and builds an inverted index based on the content of their description fields. This index enables users to search for relevant events using textual queries – much the way an Internet search engine is used. DAVVI responds with a ranked list of those events that have description fields matching the relevant keywords in the query. On the client side interface, this can for instance be presented as a list of video thumbnails, as shown in Figure 2. Using the returned search result, users can select one or more events (consisting of several smaller video segments) and compose those into a *playlist*.

Another option is to let the system automatically generate a playlist based on search result ranking and a specified length. Such playlists are used as input to the DAVVI segment delivery substrate. The listed segments are then streamed to the user for playback as a personalized video.

3. METADATA ANNOTATION

A traditional video consists of a sequence of scenes. Each scene can again contain several events. For example, a scene in a soccer game may include multiple events like “sliding tackle”, “yellow card” and “injury”. Hence, a search for any of these terms will return the same scene as one of the results. In order to allow a user to use such textual queries, a crucial task is to identify the semantics of a scene and its events. This requires very effective metadata extraction and generation tools. In this section we outline the mechanisms used in DAVVI to gather high quality metadata for video annotation and indexing.

For each video, DAVVI maintains a video description XML file. Each such description file includes general information about the video, like a unique video identity and the date. It may also contain domain specific data, like the playing teams for soccer videos. A video description file also contains a de-

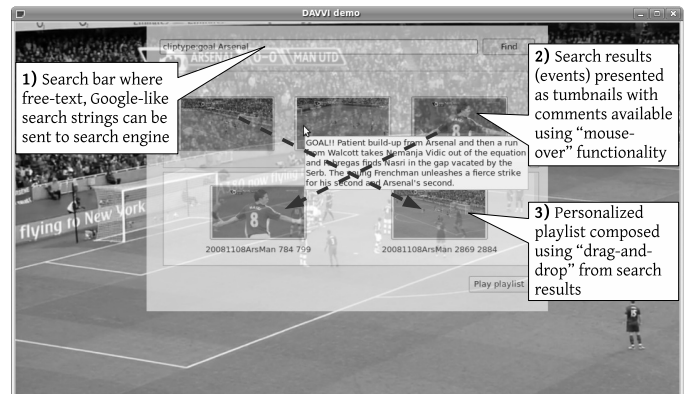


Fig. 2. Example of making a personalized video playout based on search results

scription of zero or more events that occurred within the video consisting of a time value and a description. Event descriptions are domain-specific, but typically include information like:

- Event type where common soccer events include goal, yellow card, penalty, corner, and tackle.
- Free-text annotation of the event.
- Structured XML event information.

When adding a new video to the system, the description file for that video is initialized with metadata that can be extracted from the video itself. This can for instance be textual information contained in the video's header, or features extracted by video analysis tools, which we discuss next.

3.1. Video Analysis Tools

A key challenge of video search is identifying the discrete units in the video itself for indexing. For post-production video content, a simple shot boundary segmentation is often sufficient. For more continuous content, such as user created or sports content, a more advanced event segmentation technique is required. Given an appropriate segmentation technique, each segment then needs to be annotated to support indexing and search.

When manual annotation is not possible or sufficient, automatic annotation based on visual analysis for metadata extraction directly from the video content becomes necessary. Automatic video analysis can be used to both extract metadata from the digital video and segment the video into events. Such automatic video analysis usually relies on visual analysis mechanisms to generate a segmentation or to attempt to identify and label what is happening in the video content. The first generation video analysis tools relied on modeling video with easily extractable low-level visual features such as color, texture, and edge detection. These visual features can then be used to describe the video shots in simple terms.

While such visual features have been employed successfully to support more-like-this types of video search and have been employed successfully for the TRECVID [5] collaborative video search experimentation, they are not suitable for many video search applications. The reason for this is that a significant semantic gap exists between these low-level visual features and the semantic meaning of the video content.

To retrieve events with a particular meaning, rather than only identifying colors and shapes, the semantic meaning of the video needs to be included in the metadata. To help bridge this semantic gap and make video search more user friendly, video content can be processed to seek more complex semantic concepts, such as people (faces, news readers), locations (indoor/outdoor, cityscape/landscape), objects (buildings, cars, airplanes), events (goals, explosions, car races) and production techniques such as camera motion. The output of such higher-level concept detectors can, with sufficient development and training, be integrated into content-based video retrieval systems, which can operate with varying degree of success [6].

Since the inception of this research field in the mid 90's, various higher-level techniques for automatic visual analysis of sports videos have been proposed and evaluated. Although soccer is the perhaps most widely investigated topic [7–9], other sports that have received attention are tennis [10], basketball [11, 12], baseball [13–15], rugby [16], American football [17], and Formula 1 [18]. It is generally accepted that genre-specific approaches, like soccer only or field-sport only techniques, perform better than genre-independent or more general sports event detection techniques. In either case, the approaches to event detection and summarization can be grouped according to the features they exploit into three main categories: visual analysis, audio analysis, and textual metadata analysis. There are also approaches that combine some or all of the above features into a single framework. For example, previous work by our colleagues at DCU [19] into a field-sports specific event segmentation technique utilized a Support Vector Machine (SVM) to combine both audio analysis and visual analysis to produce one combined feature-based framework for event segmentation. For every shot in the sports content, a number of feature analysis outputs were generated and these were aggregated into a shot feature vector. The number of analysis tools therefore defined the dimensionality of this feature vector. In this particular case, a five dimensional feature vector was utilized, incorporating both audio and visual analysis. The five dimension feature vector was then used as input data for an SVM event classifier which was trained to detect events in the sports content. The results of such an event segmentation technique showed 90% recall performance for sports events in soccer content with precision of 74%. There is a trade-off at work here, whereby a higher precision value results in a lower retrieval performance. There is a saturation point to be found that optimises both values.

Following a similar multi-dimensional feature vector and SVM technique, we developed and evaluated four visual and one audio feature as feature vector inputs for our detection component for important events. The visual analyses employed were: high motion activity, player zoom-in, crowd focus and goal-line. In addition to these visual analyses, the audio activity levels were also measured as an indication of the crowd and commentator excitement. It is also possible to increase precision somewhat by developing additional detectors for the noise of referee's whistle, or the noise made by the racket or bat striking the ball [20] during a tennis or baseball game. Once again, we employed an SVM as the classifier and evaluated over a number of field sports, with soccer being the best performing. Our best results show just above 85% precision at 75% level of recall. We could increase the performance of the event segmentation component somewhat by the inclusion of even more visual or audio feature detectors, especially ones that are tailored towards a given sports broadcaster.

Even though such techniques are effective, they may not suit all scenarios, for example when a very low error-rate is essential. Another problem is that the time taken to visually analyse sports content ranges from 80% of real-time for motion detection up to 300% of real-time for object tracking of players. The

31:01 GOAL - Fernando Torres
Liverpool 1 - Tottenham 0

Dirk Kuyt crosses the ball, Fernando Torres grabs a headed goal from deep inside the penalty box. Liverpool 1-0 Tottenham. Assist by Dirk Kuyt.

29:49

Foul by Didier Zokora on Steven Gerrard, free kick awarded. Free kick taken by Daniel Agger.

28:27

Free kick awarded for a foul by Steven Gerrard on Luka Modric. Gareth Bale crosses the ball in from the free kick, clearance by Jamie Carragher.

26:33

Drilled right-footed shot by Fernando Torres. Heurelho Gomes makes a save.

25:52

Shot by Jermaine Jenas from 30 yards. Jose Reina makes a save.

Fig. 3. Example live text commentary [22]

processing time typically is about 20% of real-time for excitement identification and between 100–150% of real-time for audio event detection. Clearly, this processor-intensive process introduces latency, which may not suit a live system.

3.2. External Web Sources

Due to the limitations of automatic analyzes tools, we cannot rely on these mechanisms alone. Therefore, in order to support the search functionality in DAVVI, we combine them with the text-based search detailed in Section 2. Thus, to enhance our annotations based on the large corpus of currently untapped metadata contained in the textual information made available on the web, we can use the visual analyzes tools. This is a similar approach to the one proposed in [21]. As an example, many news and broadcasting sites publish live text commentaries of important soccer events providing web pages with minute-by-minute game descriptions. For instance, VGLive, Eurosport, BBC and Setana Sports have all allocated resources to provide human-generated live commentaries for English Premier League games, like shown in Figure 3. The quality of these textual comments with second-precision timing is surprisingly well suited for our purpose, and they include all of the events, which we can not guarantee when using existing visual analysis tools. DAVVI therefore crawls such relevant web-sites, retrieves the commentaries, and annotates and builds a search index within seconds compared to the high processing requirements for automatic video analysis tools.

3.3. User Adjustments and Feedback

One additional issue stems from the fact that real-life events, like a soccer goal, are identified by time intervals (start- and stop times) and not a single moment in time. Unfortunately, the external web sources like BBC Sport do not provide information on such time-spans. As such, DAVVI’s metadata extraction engine must interpolate the event intervals, and for this first prototype, we have initially used a time-span of 30 seconds

centered on the time value listed by the external data sources. We are, however, investigating how to assign intervals based on event type and visual features, and how to use automatic tools to identify more accurate start and stop times. As a last refinement of the searchable metadata and the returned result, a user is therefore allowed to adjust the event time interval like in [23], e.g., if the returned video clip stopped in the middle of the sliding tackle. This action will have immediate effect on the played out video, and the adjustments are sent back to the system and weighted in into future search results.

4. SEARCH AND PERSONALIZED PLAYOUT

One major feature of DAVVI is that we do not only identify a rich set of events like “yellow cards”, “shot tipped over the bar”, and “sliding tackles”, i.e., events not found by automatic analysis. The users can also search for video events using textual queries and combine these arbitrarily into a personalized, content-based video playout. For example, if some user is interested in goal events where Fernando Torres is the scoring player assisted by Dirk Kuyt, then the following free-text, Google-like query can be submitted:

“goal” “Fernando Torres” “assist by Dirk Kuyt”

Additionally, the structural elements found in Figure 4 like “hometeam=Liverpool” and “how=heading” can be used to refine the query further.

Given the above query, the DAVVI search engine will return the scene at time 31:01 with the goal event shown in Figure 3. The search results are presented as for example shown in Figure 2 using a scene thumbnail and the scene description. Then, when the corresponding video of an event is to be played out, the search component forwards the correlated XML description shown in Figure 4 to the streaming component. These XML descriptions are based on aggregated metadata from, for example, Eurosport and Setana Sports and include the soccer game id and the event time attributes (time, starttime, and stoptime). This allows the streaming component to identify the corresponding video segments comprising the event which are retrieved using our segment delivery substrate described in Section 2. Using drag-and-drop functionality, the query results can also be combined arbitrarily into a playlist as shown in Figure 2, or the system can automatically generate a specified-length playout based on the search result ranking. The returned playlist is then played out as one continuous video adapting the quality to available resources, i.e., supporting on-the-fly composition of personalized, topic-based videos.

5. DISCUSSION

We have demonstrated with DAVVI that external metadata can be used to identify a large set of events within videos. However, to enable richer search in terms of the quantity and the quality of searchable items, there are still several challenges and open questions that must be addressed.

```

<doc>
  <str name="gameid">200905241ivtot</str>
  <int name="time">1861</int>
  <int name="starttime">1846</int>
  <int name="stoptime">1876</int>
  <str name="awayscore">0</str>
  <str name="homescore">1</str>
  <str name="awayteam">Tottenham</str>
  <str name="hometeam">Liverpool</str>
  <str name="scoringplayer">Fernando Torres</str>
  <str name="scoringteam">Liverpool</str>
  <str name="how">Heading</str>
  <str name="type">goal</str>
  <str name="comment">
    Dirk Kuyt crosses the ball, Fernando Torres
    grabs a headed goal from deep inside the
    penalty box. Liverpool 1-0 Tottenham. Assist
    Dirk Kuyt.</str>
</doc>

```

Fig. 4. XML sample of our search engine

A key problem is to synchronize the course-grained time information found in external data sources to the time coordinate of the corresponding video. For example, the event information that we extract from BBC Sport contains event time as minutes and seconds as shown in Figure 3. However, because the live text commentaries from BBC Sport start at kick-off, they are time-coded relative to the start of the game. Unfortunately, the videos often start a varying number of minutes before kick off. Therefore, the time of extracted events must be individually adjusted with an offset so that they match with the video. This problem is further complicated by the fact that neither game time nor video time necessarily flow at a constant rate relative to real-time. For instance, in many sports, like hockey and American football, the game time is frequently paused, or interrupted with commercial breaks and game replays. Hence, there might not be one single time offset that correctly maps all events from the external sources to our video. We are currently looking at mechanisms to find suitable time synchronization mappings by analyzing the video, like proposed in [21], to detect game start for videos that contain the official game-time clock.

So far, we have used soccer as an example, but we would like to argue that our approach can be applied to other content domains for two reasons. First, the video delivery substrate is oblivious to the video content and can be used to distribute any type of content. Second, there is a multitude of web sites providing information that can be extracted by DAVVI's metadata extraction engine. In particular, minute-by-minute live text commentaries for other sports than soccer exist, like hockey [24] and rugby [16]. More generally, our approach could be applicable to any type of videos where timed textual commentaries are provided.

As an orthogonal example to our soccer entertainment scenario, we have developed a lecture application prototype in the edutainment domain. To make the multimedia search index, we had the presentations, the transcripts, and the slide change timings available. This information was then aligned with the recorded video capture of the lecture. The search engine then

returns search results based on slide content and speech, and a video playout is generated by drag-and-dropping slide thumbnails to generate a playlist. This means that the user can on-the-fly generate a video summary from parts of various lectures talking for example about "multimedia data mining" where the newly composed presentation is displayed with the associated video scene of the person talking about the corresponding slide.

6. CONCLUSION

In this paper, we have described how we have built a content-based multimedia index using annotations derived from both video analysis and untapped textual, minute-by-minute commentaries available on the Internet. Furthermore, through a traditional search engine, we have demonstrated that relevant video events from multiple videos can be identified, retrieved and combined into one continuous video playout. In particular, we have demonstrated the usability of DAVVI in the soccer scenario, but we also conjecture a much broader applicability potential than just for the sports domain. One example is to allow students to make individual summary lectures from university courses. However, several remaining problems must be addressed to make this happen. We consider our initial use of external sport script annotations as far too application specific. Next generation DAVVI is thus aiming for a more automatic framework for content-based analysis combining the best from video analysis toolkits and metadata extraction from external sources.

The current prototype has successfully been exposed and demonstrated to a wide audience with a lot of good feedback. However, we have not yet performed extensive user evaluation tests. One of our ongoing activities is therefore to perform a more thorough evaluation, both looking at the streaming solution, video quality adaptation, search result accuracy, and user satisfaction in general.

7. REFERENCES

- [1] Dag Johansen, Håvard Johansen, Tjalve Aarflot, Joseph Hurley, Åge Kvalnes, Cathal Gurrin, Sorin Sav, Bjørn Olstad, Erik Aaberg, Tore Endestad, Haakon Riiser, Carsten Griwodz, and Pål Halvorsen, "DAVVI: A prototype for the next generation multimedia entertainment platform (demo)," in *Proc. of the ACM International Multimedia Conference (ACM MM)*, Oct. 2009, pp. 989–990.
- [2] Pengpeng Ni, Alexander Eichhorn, Carsten Griwodz, and Pål Halvorsen, "Fine-grained scalable streaming from coarse-grained videos," in *Proc. of the International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, June 2009.
- [3] Haakon Riiser, Pål Halvorsen, Carsten Griwodz, and Dag Johansen, "Low overhead container format for adaptive streaming," in *Proc. of the Multimedia Systems Conference (MMsys)*, Feb. 2010.

- [4] Mohammed A. Refaey, Wael Abd-Almageed, and Larry S. Davis, "A logic framework for sports video summarization using text-based semantic annotation," in *Proc. of the International Workshop on Semantic Media Adaptation and Personalization (SMAP)*, Washington, DC, USA, 2008, pp. 69–75.
- [5] Alan F. Smeaton, Paul Over, and Wessel Kraaij, "Evaluation campaigns and TRECVID," in *Proc. of the ACM international workshop on Multimedia information retrieval (MIR)*, 2006, pp. 321–330.
- [6] Cathal Gurrin, "Content-based video retrieval," in *Encyclopedia of Database Systems*, pp. 466–473. Springer US, 2009.
- [7] Dennis Yow, Boon lock Yeo, Minerva Yeung, and Bede Liu, "Analysis and presentation of soccer highlights from digital video," in *Proc. of the Asian Conference on Computer Vision*, 1995, pp. 499–503.
- [8] Okihisa Utsumi Koichi, Koichi Miura, Ichiro Ide, Shuichi Sakai, and Hidehiko Tanaka, "An object detection method for describing soccer games from video," in *Proc. of the IEEE International Conference on Multimedia and Expo (ICME)*, 2002, pp. 45–48.
- [9] Ahmet Ekin and A. Murat Tekalp, "Automatic soccer video analysis and summarization," *IEEE Trans. on Image Processing*, vol. 12, pp. 796–807, 2003.
- [10] Rozenn Dahyot Anil, Anil Kokaram, Niall Rea, and Hugh Denman, "Joint audio visual retrieval for tennis broadcasts," in *Proc. of the IEEE Conf. Acoustics, Speech, and Signal Processing, Hong Kong*, 2003, pp. 561–564.
- [11] Wensheng Zhou, Asha Vellaikal, and C. C. Jay Kuo, "Rule-based video classification system for basketball video indexing," in *Proc. of the ACM Multimedia workshops*, 2000, pp. 213–216.
- [12] D. D. Saur, Y.-P. Tan, S. R. Kulkarni, and P. J Ramadge, "Automated analysis and annotation of basketball video," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Jan. 1997, vol. 3022, pp. 176–187.
- [13] Yong Rui, Anoop Gupta, and Alex Acero, "Automatically extracting highlights for tv baseball programs," in *MULTIMEDIA '00: Proc. of the eighth ACM international conference on Multimedia*, New York, NY, USA, 2000, pp. 105–115, ACM.
- [14] Peng Chang, Mei Han, and Yihong Gong, "Extract highlights from baseball game video with hidden markov models," in *Proc. of the IEEE ICIP 2002*, 2002, pp. 609–612.
- [15] T. Kawashima, K. Tateyama, T. Iijima, and Y. Aoki, "Indexing of baseball telecast for content-based video retrieval," in *ICIP 98 : Proc. of Image Processing*, 1998, pp. 871–875.
- [16] RTÉ Commercial Enterprises Limited, "Rugby news," Web page, <http://www.rte.ie/sport/rugby>, 2009.
- [17] B. Li and M. I. Sezan, "Event detection and summarization in American football broadcast video," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Dec. 2001, vol. 4676, pp. 202–213.
- [18] Milan Petkovic Vojkan, Milan Petkovic, Vojkan Mihajlovic, Willem Jonker, and S. Djordjevic-kajan, "Multimodal extraction of highlights from tv formula 1 programs," in *Proc. of the IEEE International Conference on Multimedia and Expo (ICME)*, 2002, pp. 817–820.
- [19] Yi Ding and Guoliang Fan, "Event detection in sports video based on generative-discriminative models," in *EiMM '09: Proc. of the 1st ACM international workshop on Events in multimedia*, New York, NY, USA, 2009, pp. 17–24, ACM.
- [20] Chunxi Liu, Qingming Huang, Shuqiang Jiang, Liyuan Xing, Qixiang Ye, and Wen Gao, "A framework for flexible summarization of racquet sports video using multiple modalities," *Comput. Vis. Image Underst.*, vol. 113, no. 3, pp. 415–424, 2009.
- [21] Changsheng Xu, Jinjun Wang, Kongwah Wan, Yiqun Li, and Lingyu Duan, "Live sports event detection based on broadcast video and web-casting text," in *Proc. of the ACM International Multimedia Conference (ACM MM)*, Oct. 2006, pp. 221–230.
- [22] British Broadcasting Corporation (BBC), "Premier league live scores," Web page, http://news.bbc.co.uk/sport2/hi/football/eng_prem/live_scores, 2009.
- [23] Andrew Brampton, Andrew MacQuire, Michael Fry, Idris A. Rai, Nicholas J. P. Race, and Laurent Mathy, "Characterising and exploiting workloads of highly interactive video-on-demand," *ACM/Springer Multimedia Systems*, vol. 15, no. 1, pp. 3–17, Feb. 2009.
- [24] Lars Wiklund, Thomas Mattsson, and Joakim Svensson, "HockeyExpressen," Web page, AB Kvällstidningen Expressen, <http://hockey.expressen.se/>, Apr. 2009.