




US 20130293776A1

(19) **United States**(12) **Patent Application Publication**  
**Shih-Fu et al.**(10) **Pub. No.: US 2013/0293776 A1**(43) **Pub. Date: Nov. 7, 2013**(54) **SYSTEM AND METHOD FOR EXTRACTING  
TEXT CAPTIONS FROM VIDEO AND  
GENERATING VIDEO SUMMARIES****Publication Classification**(51) **Int. Cl.**  
**H04N 7/088** (2006.01)  
(52) **U.S. Cl.**  
CPC ..... **H04N 7/0882** (2013.01)  
USPC ..... **348/468**(71) Applicant: **The Trustees of Columbia University,**  
New York, NY (US)(72) Inventors: **Chang Shih-Fu**, New York, NY (US);  
**Dongqing Zhang**, Burbank, CA (US)(21) Appl. No.: **13/935,183**(22) Filed: **Jul. 3, 2013****Related U.S. Application Data**(60) Division of application No. 11/960,424, filed on Dec.  
19, 2007, now Pat. No. 8,488,682, which is a continu-  
ation of application No. 10/494,739, filed on May 8,  
2004, now Pat. No. 7,339,992, filed as application No.  
PCT/US2002/039247 on Dec. 6, 2002.(60) Provisional application No. 60/337,911, filed on Dec.  
6, 2001.(57) **ABSTRACT**

Caption boxes which are embedded in video content can be located and the text within the caption boxes decoded. Real time processing is enhanced by locating caption box regions in the compressed video domain and performing pixel based processing operations within the region of the video frame in which a caption box is located. The captions boxes are further refined by identifying word regions within the caption boxes and then applying character and word recognition processing to the identified word regions. Domain based models are used to improve text recognition results. The extracted caption box text can be used to detect events of interest in the video content and a semantic model applied to extract a segment of video of the event of interest.






FIG.1

FIG. 2

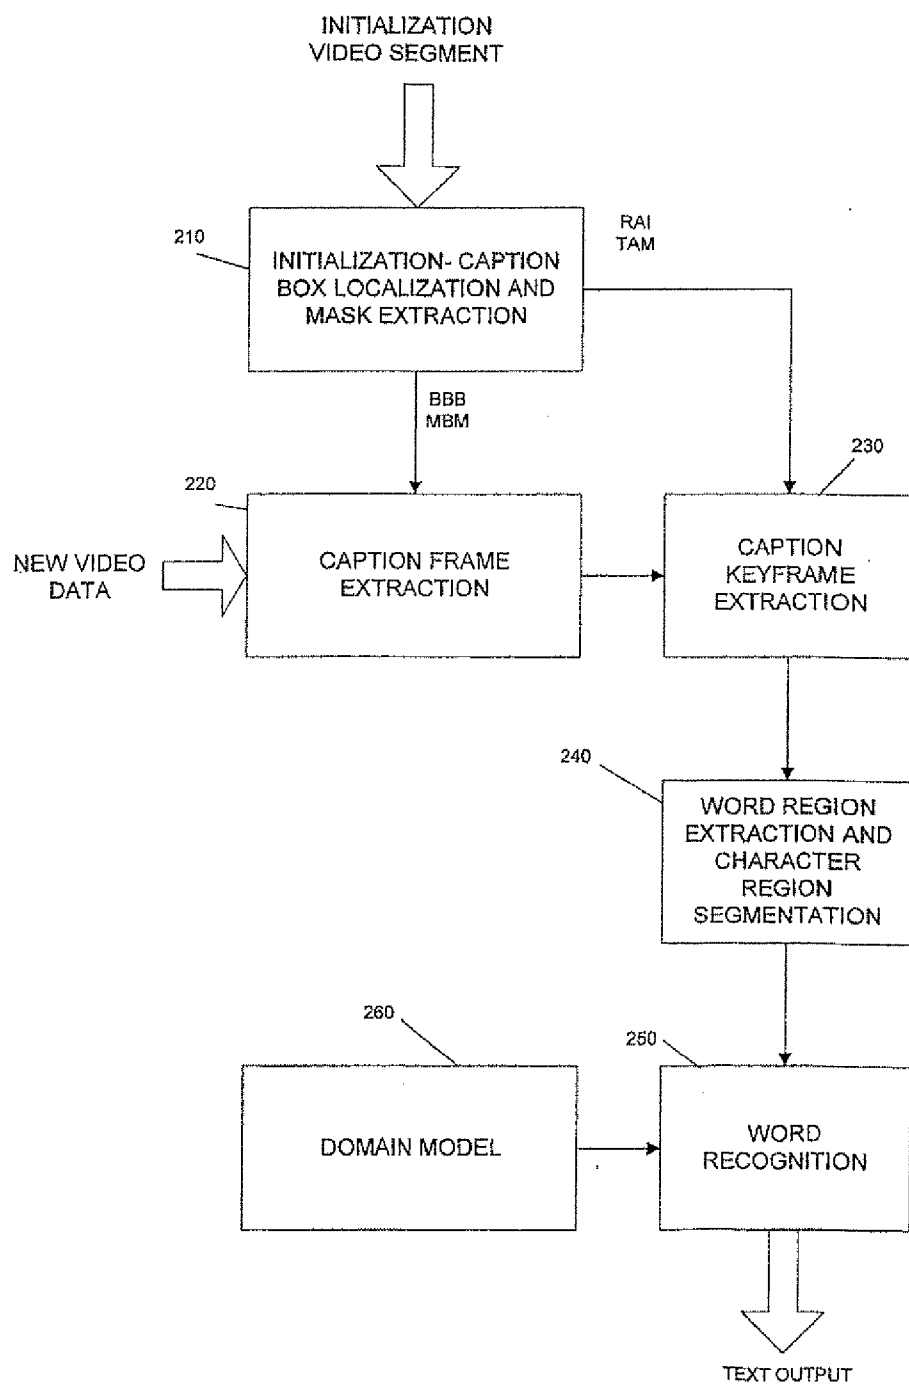
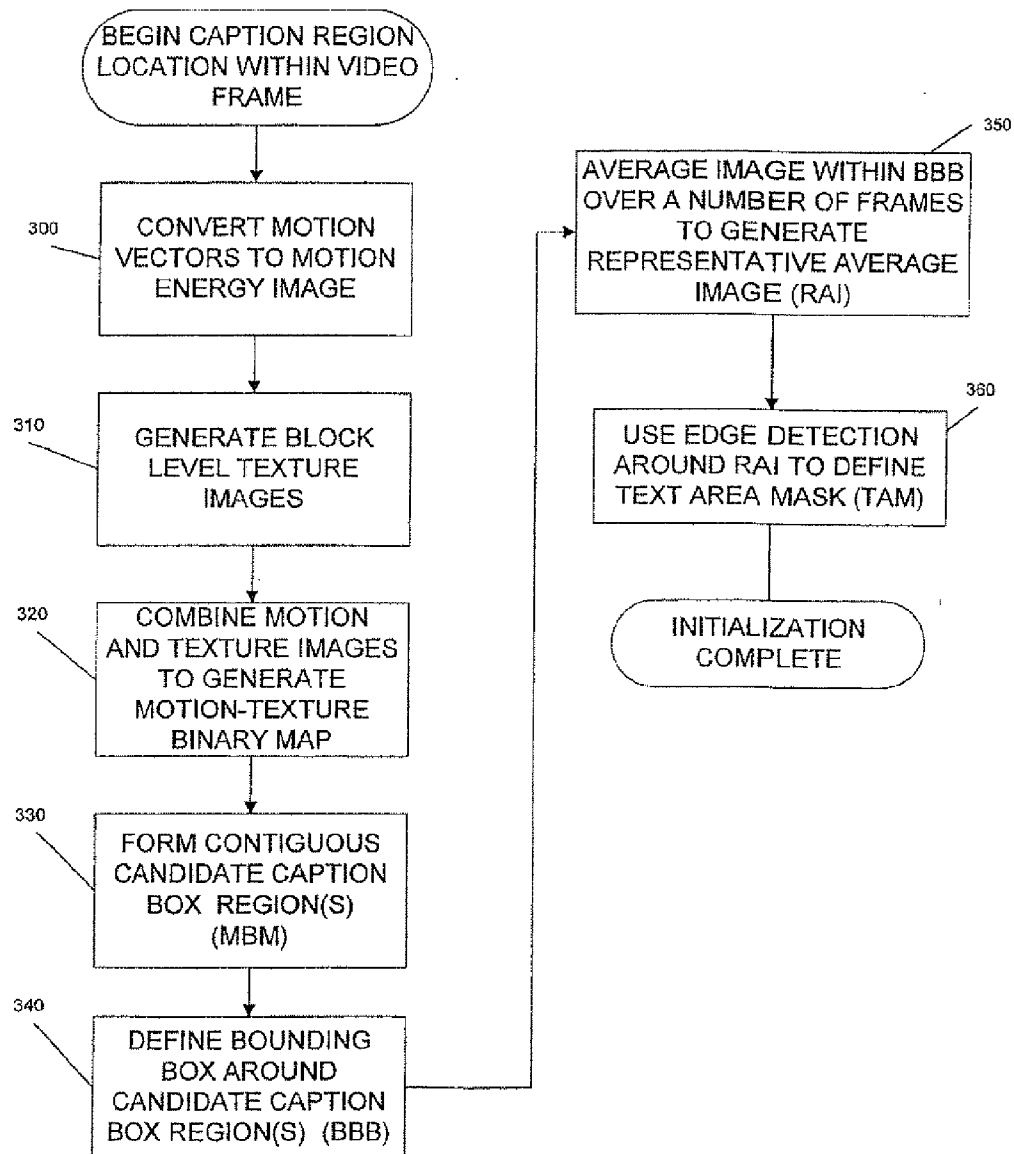


FIG. 3






FIG. 4




FIG. 5


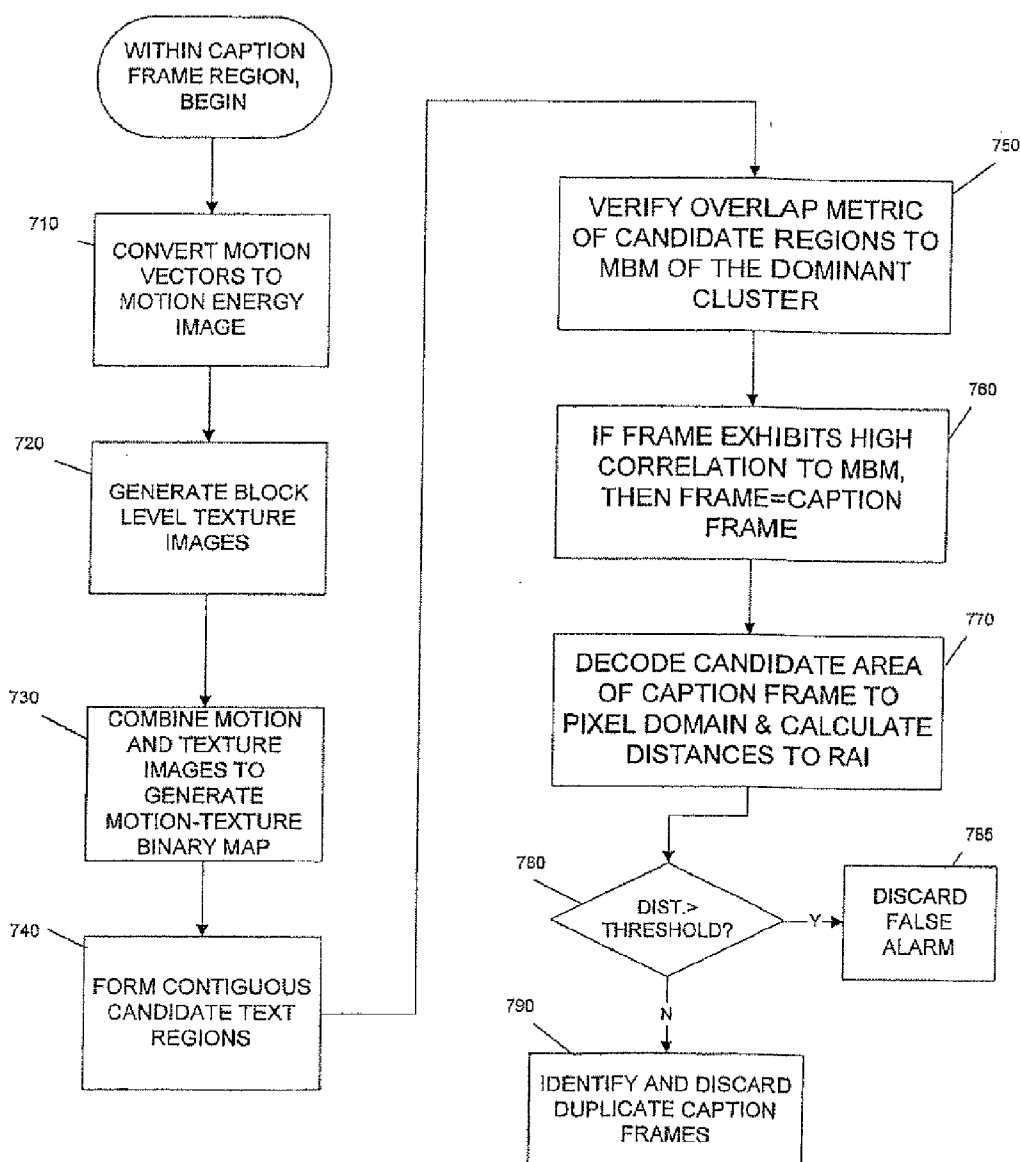


FIG. 6

FIG. 7






FIG. 8A




FIG. 8B





FIG. 8C

FIG. 9








FIG. 10 PRIOR ART




FIG. 11

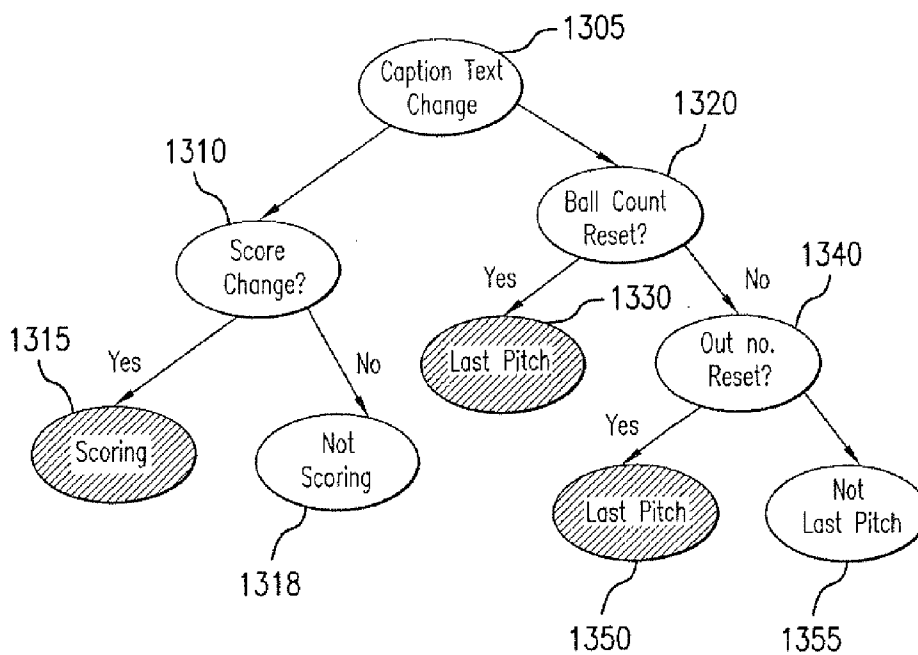


FIG. 13




FIG. 12




FIG. 14

## SYSTEM AND METHOD FOR EXTRACTING TEXT CAPTIONS FROM VIDEO AND GENERATING VIDEO SUMMARIES

### FIELD OF THE INVENTION

**[0001]** The present invention relates generally to text recognition, and more particularly relates to the detection and decoding of caption regions embedded in video content and using the extracted text to generate video summaries.

### BACKGROUND OF THE INVENTION

**[0002]** There exists a substantial volume of video and multimedia content which is available both online, such as via the Internet, and offline, such as in libraries. In such video and multimedia content, it is common for a text caption box to be embedded in the video to provide further information about the video content. For example, as illustrated in FIG. 10, a video recording of a baseball game typically includes a caption box 1010 which displays game statistics such as the score, inning, ball/strike count, number of outs, etc. The detection and recognition of the text captions embedded in the video frames can be an important component for video summarization, retrieval, storage and indexing. For example, by extracting a short video segment preceding certain changes in the text of the baseball caption box, such as score or number of outs, a "highlight" summary can be automatically generated.

**[0003]** Text recognition in video has been the subject of current research. For example, the article "Video OCR: Indexing Digital News Libraries by Recognition of Superimposed Captions," by T. Sato, et al., Multimedia Systems, 7:385-394, 1999 discloses a system for detecting and recognizing text in news video. This system is described as using a line filter to enhance the text characters and a projection histogram to segment the characters. A dynamic programming algorithm is used to combine the segmentation and recognition processes to reduce the false alarms of character segmentation.

**[0004]** Past approaches to text detection in video do not adequately account for disturbances in the background areas. As a result, previous approaches are often sensitive to cluttered backgrounds, which diminish text recognition accuracy. Therefore, there remains a need for improved methods of extracting text embedded in video content. There also remains a need to improve automatic video summary generation methods using text which is extracted from the video content.

### SUMMARY OF THE INVENTION

**[0005]** It is an object of the present invention to provide a system and method for location and recognition of text embedded within video content.

**[0006]** It is a further object of the present invention to provide a method of locating a caption box within video content and recognizing the text within the caption box.

**[0007]** It is another object of the present invention to provide a system and method for identifying a caption box in video content in the sports domain and detecting changes in the game state based on the text in the caption box.

**[0008]** It is yet another object of the present invention to provide a method of generating a summary of video content by detecting a caption box and selecting video segments for the summary based on changes within the caption box.

**[0009]** In accordance with the present invention, a method of decoding a caption box in video content is provided. In the method, the expected location of a caption box in a frame of the video content is determined. At least one caption box mask within the expected location is also determined. A caption box mask is applied to frames of the video content and those frames exhibiting a substantial correlation to the caption box mask within the expected caption box location are identified as caption frames. For at least a portion of the caption frames, word regions within the confines of the expected location are identified and within each word region, text characters are identified. The text characters in the word region are compared against a domain specific model to enhance word recognition.

**[0010]** In the present method, determining an expected location of a caption box can include evaluating motion features and texture features of the video frame in the compressed domain and identifying regions having low motion features and high texture features as candidate caption box regions.

**[0011]** To enhance processing efficiency, it is desirable to remove duplicate caption frames from word region processing. Therefore, the method can further include evaluating the identified caption frames, within the caption box location, for changes in content; and removing caption frames from word region processing which do not exhibit a change in content. Alternatively, a subset of the caption frames can be selected for word region processing by selecting caption frames spaced over a predetermined time period.

**[0012]** In one embodiment, the operation of identifying text characters includes generating a vertical projection profile for each word region and identifying local inflection points, such as minima, in the vertical projection profile. Character regions can then be defined by selecting those minima which are below a threshold value as the position of character boundaries in the word region. A character recognition algorithm is then used to evaluate the defined character regions.

**[0013]** Also in accordance with the present invention is a method of generating an event based summary of video content which includes caption boxes embedded therein. The summarization method begins by extracting caption boxes from at least a portion of the frames of the video content and identifying changes in the content of the extracted caption boxes which are indicative of an event of interest. For each identified change in the content of the caption box, a semantic model is applied to select a portion of the video content, preceding the change in the content of the extracted caption box, which includes the event of interest.

**[0014]** The above described method of caption box extraction and decoding can be used in the summarization method to identify changes in the content of the extracted caption boxes

**[0015]** In one embodiment of the summarization method, the video content is of a baseball game. In this domain, the semantic model can identify the portion of the video content of the event of interest as residing between a pitching event and a non-active view. In this regard, the pitching event can be identified using color model matching and object layout verification, such as the typical arrangement of the pitcher, batter and field. Non-active view frames, which generally include views of the audience or non-active players, can be identified by a reduction in the number of green pixels as compared to a preceding frame as well as a decrease in motion intensity.

## BRIEF DESCRIPTION OF THE DRAWING

[0016] Further objects, features and advantages of the invention will become apparent from the following detailed description taken in conjunction with the accompanying figures showing illustrative embodiments of the invention, in which:

[0017] FIG. 1 is a pictorial representation of a frame of sports video content which includes a text caption box;

[0018] FIG. 2 is a simplified operational block diagram of an embodiment of the present system for locating a caption box in video content and extracting text therefrom;

[0019] FIG. 3 is a simplified flow diagram illustrating the process of determining the caption box location within the video frames;

[0020] FIG. 4 is a pictorial representation of a block level mask used to identify frames having caption boxes located therein;

[0021] FIG. 5 is a pictorial representation of an exemplary average image of a caption box region;

[0022] FIG. 6 is a pictorial representation of a text area mask derived from the average image of FIG. 5.

[0023] FIG. 7 is a flow diagram further illustrating a process of identifying caption frames in video content;

[0024] FIG. 8A is a pictorial representation of a word region extracted in accordance with the present invention;

[0025] FIG. 8B is a pictorial representation of the word region of FIG. 8A following intensity segmentation;

[0026] FIG. 8C is a graph illustrating a vertical intensity profile versus pixel position for the word region of FIG. 8A;

[0027] FIG. 9 is a transition graph illustrating expected state transitions in a ball-strike count encountered in video content of a baseball game;

[0028] FIG. 10 is a pictorial diagram illustrating a frame of video with a caption box embedded therein, which is known in the prior art.

[0029] FIG. 11 is a temporal flow diagram illustrating a sequence of semantic events which occur in baseball video content;

[0030] FIG. 12 is a pictorial representation of a caption box and further illustrating word region types in a baseball video domain;

[0031] FIG. 13 is a flow diagram illustrating the process of identifying scoring events and last pitch events for generating a summary of a baseball video content from extracted caption box text; and

[0032] FIG. 14 is a pictorial representation of an exemplary computer display of a video summary browser generated in accordance with the present invention.

[0033] Throughout the figures, the same reference numerals and characters, unless otherwise stated, are used to denote like features, elements, components or portions of the illustrated embodiments. Moreover, while the subject invention will now be described in detail with reference to the figures, it is done so in connection with the illustrative embodiments. It is intended that changes and modifications can be made to the described embodiments without departing from the true scope and spirit of the subject invention as defined by the appended claims.

## DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0034] The present invention is directed to systems and methods of locating a caption box within video content and

then applying methods of optical character recognition to extract the text from the caption box. Domain knowledge can be applied to the extracted text to enhance the character recognition results. By identifying changes in the detected text which represent events of interest, the extracted caption box text can be used to form a text based and/or a video based summary of the video content.

[0035] Referring to FIG. 1, video content can be represented as a sequence of discrete frames **100** through **100-N** which are displayed over time. The video content can be provided in compressed form, such as using MPEG compression, as well as in high resolution pixel based form. Those frames which include a caption box displayed therein are referred to herein as caption frames. Those caption frames where the content within the caption box undergoes a cognizable change from the previous caption frame are referred to as keyframes.

[0036] An overview of the operation of the present invention is provided in connection with FIG. 1, which is a pictorial representation of a segment of video content having a number of frames **100** sequenced in time (**100-1** through **100-N**). In frame **100**, a caption box region, referred to as a block level bounding box **105** is identified. As is described in greater detail below, the block level bounding box **105** is located during an initialization procedure and is used to confine processing of subsequent frames **100-1** through **100-N** to this portion of the video frame. Within the bounding box **105**, one or more templates will be defined to identify and process the caption box image **110**. Within the caption box image **110**, word regions **115** will be identified. The word regions **115** will be further processed to define character boundaries which define character regions **120** within the word regions. The pixels within the character regions **120** will preferably be subjected to character recognition processing and the word regions will be subjected to a domain specific model for word level verification, word region type identification and correction.

[0037] FIG. 2 is a simplified operational block diagram of a system for locating a caption bounding box **105** in the video frame **100**, identifying frames having a caption box therein, and extracting text from at least a portion of the video frames which include a caption box. It is assumed that the caption box will be consistently placed in the same region of the video frame throughout a particular piece of video content. However, because there is no predefined standard for the placement of a caption box within the frame, different video sources may place a text caption box in various locations in the video frame and may place the text within the caption box differently. Therefore, the present method begins with an initialization process which determines the location of candidate caption box bearing areas and defines masks for processing the caption boxes (step **210**).

[0038] The present method for caption bounding box detection relies on two image characteristics of a caption box. First, the position of a caption box from frame to frame, when it is displayed, will remain essentially constant. Therefore, motion vectors for the DCT blocks in the compressed domain (MPEG) for caption box regions of the video frame will be small. Second, text is generally displayed with relatively high contrast against a background. Thus, if a text region is mapped with pixel location in the X-Y plane and pixel intensity in the Z-axis, text regions generally result in rapid and extreme changes along the Z-axis, and therefore can be considered as highly textured regions. Thus, regions within a

video frame which exhibit low motion and high texture will be candidates for caption box regions.

**[0039]** FIG. 3 is a flow diagram which further illustrates the operation of the initialization processing block 210 which includes determining the location of a caption bounding box 115 in a video frame 100. Referring to FIG. 3, video content, preferably in the compressed domain, is used to identify a bounding box 115 around candidate caption box regions. The most common compressed domain video format follows the MPEG compression standard, such as MPEG-1, although other compression schemes can be used in the practice of the present invention. In MPEG-1 compressed video, pixels of the uncompressed video are represented by blocks, which represent a discrete portion of the uncompressed video (such as 1 block=8 pixel×8 pixel region). The blocks are farther grouped as macro-blocks which contain four blocks (2 blocks×2 blocks). The macro-blocks include motion vectors. Motion vectors, macro-blocks and blocks are set forth in the MPEG-1 standard and are well understood in the art.

**[0040]** In step 300, the motion vectors from the compressed MPEG video data are converted into a macro-block level motion energy image for the video frame 100. After forming the motion energy image, the motion energy image can be upsampled in width and height by the corresponding number of blocks in the macro-block width and height respectively. This upscaling operation translates the macro-block level motion energy image into the same size as a block level image. As noted above, caption box regions within the frame are generally static. Therefore regions exhibiting low motion energy are indicative of caption box regions. Using thresholding, the block level motion energy image can be converted to a binary motion image that indicates which regions in the frame have motion values below the threshold.

**[0041]** In addition to motion, texture is also an indicator of text regions. The discrete cosine transform (DCT) coefficients on I frames in the compressed domain MPEG video can be used to extract the texture features of the video frame (step 310). A texture energy image is preferably generated on a block level of the MPEG video. The texture features indicate how rapidly the contrast is changing in that block over successive frames. As noted above, text regions are generally characterized as highly textured. Using thresholding, the texture energy images can be binarized into binary texture images which distinguish between high texture regions and lower texture regions. The result is a binary texture image.

**[0042]** In step 320 the binary motion energy image from step 300 and the binary texture image from step 310 are combined to form a joint motion-texture binary map. The images can be combined over neighboring I and B/P frames in the MPEG video data. The binary images are preferably combined using a Boolean logic AND function which removes regions that do not exhibit both low motion and high texture. In the resulting joint motion-texture map, the binary value "1" corresponds to blocks in the compressed domain which are candidate text blocks, i.e., blocks exhibiting substantially no motion and high texture.

**[0043]** The candidate text blocks in the compressed video frame are then evaluated to form contiguous candidate caption box regions (step 330). Many known forms of connected component analysis can be applied to form such regions. For example, a candidate text block can be selected in the motion-texture map as a seed block. The neighbors of the seed block can then be evaluated to determine if they represent candidate text blocks or boundary blocks. Neighboring candidate text

blocks are added to the region whereas neighboring blocks which are not candidate text blocks represent the boundary of the region which terminate region growing in that direction. The region growing process continues until each neighbor has been evaluated and either added to the region or boundary.

**[0044]** The process of forming contiguous regions is repeated throughout the video frame by selecting additional seed blocks that are not already part of a region and performing the region growing analysis until all candidate text blocks in the frame have been processed. Morphological filtering can be applied to the regions in order to remove spurious areas from the candidate regions. It will be appreciated that this form of seed based region growing is but one example of connected component analysis which can be used to identify groups of contiguous regions of candidate text blocks.

**[0045]** The candidate text regions will generally include a number of regions which do not represent likely caption box locations. In order to reduce the number of false alarms resulting from these candidates, the candidate text regions can be evaluated over a number of frames over a time window and an incremental clustering method can be applied. For example, the candidate regions in each frame can be mapped to region clusters based on the following area overlap metric:

$$S_r(R_1, R_2) = 1 - \{a(R_1 - R_1 \cap R_2) + a(R_2 - R_2 \cap R_1)\} / \{a(R_1) + a(R_2)\} \quad \text{EQ 1}$$

where  $R_1$ ,  $R_2$  are the two regions in consecutive frames and  $a(R)$  is the area of the region  $R$ . A new region is mapped to an existing cluster if the above metric is less than a threshold value, otherwise a new cluster is formed. The clustering process stops when a dominant cluster is identified. A dominant cluster is one in which at least a predetermined minimum number (such as 40%) of the frames in a continuous sliding window (e.g. 30 seconds) are mapped to the cluster.

**[0046]** The dominant cluster is used to generate a block level mask in the compressed video domain which is referred to as a Median Binary Mask (MBM) 410, such as is illustrated in FIG. 4. The MBM 410 is generated by taking the median value of the member binary images forming the dominant cluster. A rectangular bounding box around the MBM, referred to as a Block-level Bounding Box (BBB), is also defined (step 340). An example of a BBB is illustrated by bounding box 105 in FIG. 1 as well as by the outer rectangular extent 400 of the MBM 410. The MBM 410 and BBB 400 are outputs of the caption box localization process 210 which are passed to a caption frame extraction processing block 220.

**[0047]** The BBB 105 is used to constrain processing operations in subsequent video frames (100-1 to N) to the region of the frame in which the caption box is expected to reside. The MBM is applied within this region to determine if the caption box is present in a current frame. Those frames in which a caption box is present are referred to as caption frames. By limiting post-initialization processing to the region within the BBB, processing efficiency is improved.

**[0048]** Returning to FIG. 3, in addition to compressed domain block level operations which generate the MBM and BBB, the caption box localization and mask extraction processing block 210 also includes pixel domain processing within the BBB. In step 350, an averaging operation is applied to the pixels within the image constrained by the BBB, referred to as a caption image, to obtain a Representative Average Image (RAI). FIG. 5 is a pictorial example of an RAI.

**[0049]** Since text pixels and text background remain generally static from frame to frame while the image pixels

outside a caption box generally vary over time, the averaging operation tends to reinforce the caption box pixels while smoothing out temporal variations in the surrounding image pixels. In step 360, the RAI is subjected to edge detection in order to extract the outer contour of the text area and define a Text Area Mask (TAM) 610, an example of which is illustrated in FIG. 6. Various forms of edge detection which are known to those skilled in the art of image processing can be used to define the TAM 610. The RAI and TAM are outputs from the initialization process that are provided from the caption box localization and mask extraction processing block 210 to a caption keyframe extraction processing block 230.

[0050] The initialization processing of the caption box localization and mask extraction block 210 generally takes on the order of 30-60 seconds of video content, after which the outputs BBM, MBM, RAI, and TAM are used in a continuous process to identify caption frames and extract the caption box text from the video content data which is input into the caption frame extraction processing block 220.

[0051] Returning to FIG. 2, after initialization processing is complete in block 210, the video content is applied to the caption frame extraction processing block 220 and is analyzed to detect frames that are caption frames. The caption frame extraction processing block 220 performs processing of the video content in the compressed domain in a portion of the video frame constrained by the BBB. Constraining the caption frame extraction processing operations allows for efficient real time processing since only a relatively small portion of the video frame is subjected to processing.

[0052] The operation of the caption extraction processing block 220 is further described in connection with the flow chart of FIG. 7. Processing operations 710 through 760 are preferably performed in the compressed video domain. Processing is limited in each frame to the region defined by the BBB. Steps 710 through 740 are analogous to steps 300 through 330 described in connection with FIG. 3. In step 710, the motion vectors of the compressed video blocks within the BBB are used to form a binary motion energy image. In addition, a binary texture image is generated for the area within the BBB to identify blocks exhibiting high texture levels which are indicative of text (step 720). The binary texture image and the motion energy image are combined in step 730 to generate a motion-texture binary map for the region of the current frame within the BBB. Contiguous regions within the motion-texture binary map are formed using region growing or other connected component analysis techniques to identify a dominant cluster in step 740. The overlap metric of Equation 1 can then be used to compare the dominant cluster of the current frame to the MBM generated during initialization (step 750). If the overlap metric indicates that there is a high correlation between the dominant cluster in the current frame and that of the MBM, then the frame is identified as a caption frame in step 760.

[0053] Word region detection and character recognition are processing intense operations. Thus it is desirable to minimize not only the area within a frame which is subjected to such processing but also to minimize the number of frames which are subjected to such processing.

[0054] Returning to FIG. 2, after caption frames are identified, it is often desirable to further identify those caption frames in which the content of the caption box changes, since these frames will generally provide the most valuable information regarding state changes in the video content. Such

caption frames are referred to as keyframes. Keyframe extraction from the set of caption frames identified in step 220 is performed in the caption keyframe extraction block 230.

[0055] Keyframe extraction is further illustrated in connection with FIG. 7. For a frame identified as a caption frame in step 760, the region within the BBB is decoded from the compressed domain into the pixel domain (step 770). The pixel domain caption image data is then compared against the RAI to eliminate false alarms. For example, a pixel-wise Euclidean distance in the RGB space between the caption image of the current frame and the RAI can be computed and compared against a predetermined threshold value (step 770). If in step 780 the distance exceeds the threshold value, the frame can be discarded as a false alarm (step 785). Duplicate caption frames are identified and discarded in step 790 such that the caption frames which remain are a set of caption frames in which the content changes.

[0056] Duplicate caption frames can be identified by comparing the caption image of a current caption frame against the caption image of the preceding caption frame and determining the number of pixels which have changed from high intensity values to low intensity pixel values. Alternatively, comparing the caption images in consecutive frames and calculating a pixel-wise Euclidean distance in successive caption frames can be used to identify and eliminate caption frames having duplicate content. Those frames which remain will represent caption keyframes.

[0057] As an alternative, it has been found that keyframe identification processing can be eliminated and that satisfactory results obtained by simply selecting every nth caption frame for subsequent word region extraction processing. The time between selected nth caption frames will depend on the content and the expected rate of change in the caption frames for the particular domain.

[0058] Returning once more to FIG. 2, after keyframes are identified or a subset of the caption frames are selected in processing block 230, word region extraction processing and character region segmentation processing block 240 is applied to identify word regions 115 within the caption box (FIG. 1). Word region extraction is performed on the selected caption frames or keyframes in the pixel domain in a region confined to the TAM. FIG. 8A is a pictorial representation of a word region in the pixel domain which includes the characters NY. To identify word regions, the grayscale values of the caption image are converted to a binary image by applying a suitable thresholding operation. Pixels having an intensity value higher than the threshold are considered to be text pixels. Connected component analysis, such as region growing, can be applied to the text pixels in the binary image to form contiguous candidate word regions. Size constraints can then be applied to the candidate regions to eliminate those regions which are either too large or too small to be considered valid word regions. Conditions can also be placed on the aspect ratio of valid word regions.

[0059] It has been found that graphic objects within a caption box in the proximity of the text, such as lines, boxes, logos and the like, may interfere with spatial segmentation used to identify word regions and can result in improper identification, or non-identification, of word regions in the vicinity of such objects. To correct for this potential source of error, a temporal filtering operation can be used. In one form of temporal filtering, a temporal variance of the intensity of each pixel in the word regions is calculated using buffered caption images over a predetermined time window. The



resulting variance map is subjected to thresholding and region growing to identify those word regions which exhibit changes over time (such as the score or ball-strike count). The temporal filtering can remove regions with static values by thresholding temporal variances of the image. However, as certain desired fields will remain static, such as team names, temporal filtering should be used with constraints that do not negatively impact such desired regions. Also, since temporal filtering requires buffering a large number of caption images, it may be too processing intense for many real time applications. In these cases, careful selection of the binarization threshold value used to identify the word regions can substantially eliminate problems arising from static graphic elements in the caption box.

**[0060]** Referring to FIG. 1, within the word regions **115**, character segmentation is applied to identify character boundaries which define character regions **120** within the word region **115**. In this regard, each word region **115** is traversed in a pixel wise fashion across the width of the word region. At each pixel (column) location along the width of the word region, a vertical projection profile value is determined by summing the pixel values at that pixel location along the word region height.

**[0061]** FIG. 8C illustrates an example of a graph of vertical projection profile values versus pixel position for the word region of FIG. 8A. The graph exhibits local minima **810** and **820**. Local minimum **810**, which has a value higher than threshold value **830**, represents the diagonal stroke connecting the two vertical strokes in the letter N. Assuming that the text characters are presented as light text against a darker background (such as white text on a black background), local minima in the vertical projection profile graph, which are below a threshold value illustrated by line **830**, such as minimum **820**, represent dark regions between characters. Such local minima can be used to partition the word region into individual character regions. Thus, local minimum **820** establishes the location of character partition **840** between the characters N and Y in the word region illustrated in FIG. 8A. It will be appreciated that in the event of dark characters against a light background, local maxima of the vertical projection profile graph would provide the same character partition information as is provided by the local minima in FIG. 8C. Local minima and local maxima can collectively be referred to as local inflection points.

**[0062]** In addition to evaluating the vertical projection profile of the word regions, a pixel intensity histogram can be calculated within the word regions and this histogram used to classify the pixels within the word regions. FIG. 8B illustrates the result of a three level histogram thresholding applied to the image of FIG. 8A. Those pixels exhibiting values in the highest intensity portion of the histogram can be considered character pixels **850** whereas the pixels exhibiting values in the lowest portion of the histogram can be considered background pixels **870**. Pixels exhibiting an intermediate value can be considered as "other" pixels **860**.

**[0063]** For each of the character regions that are identified, the character type pixels **850** are subjected to a character recognition processing operation. There are many known methods for performing character recognition which are suitable for use in the present system. A preferred method uses Zernike moments, such as is described in "Feature Extraction Methods for Character Recognition—A Survey," by Trier, et al., Pattern Recognition, vol. 29, pp. 641-662, 1996, the disclosure of which is hereby incorporated by reference. Zernike

moments are calculated based on a set of complex polynomials that form a complete orthogonal set inside the unit circle. The Zernike moments are the projection of the image onto these complex bases.

**[0064]** Prior to calculating the Zernike moments, it is preferable to convert the character image into three binary images using the threshold values established in connection with the intensity histogram for each word region which identifies pixels as character, background or other pixels. The first binary image distinguishes background pixels vs. other pixels. The second binary image distinguishes other pixels vs. character pixels. The third binary image is formed using the average value of the two thresholding values. Zernike feature vectors are then calculated for each of the three images. The use of three binary images enhances the robustness of the character recognition process with respect to variations in character line width and font type.

**[0065]** For each character region **120**, the image of the character is projected to the complex bases to form a character feature vector. To calculate Zernike moments, the coordinate of each character pixel is normalized into the interior of the unit circle, i.e.,  $x^2 + y^2 \leq 1$ . Since the output of the Zernike filter is a complex vector, each component of the vector includes a real and an imaginary component. The real and imaginary components can be converted to the magnitude and phase of the Zernike moments. It has been found that using nineteen (19) as the maximum order of Zernike moment and 220 as the dimension of the character feature vector provides satisfactory results.

**[0066]** Accuracy in the character recognition process can be improved by employing one or more domain models to verify conventional character recognition processing. Domain models can include specialized dictionaries and state transition models which represent expected state changes in the text regions. In addition, various word regions can be identified by a semantic word type, such as team name, score, inning and the like, and specific models applied to the individual word regions based on the word type. For example, in a baseball video, the number of team names or abbreviations is known and can be used to identify and verify a team name word region.

**[0067]** For each entry in the domain specific library, the Zernike feature vectors for each character in the word can be concatenated to generate word-level Zernike feature vectors for the dictionary entries. During word region recognition, for each word region identified in the caption box, the Zernike character feature vectors can be concatenated to generate word-level feature vectors which can be tested against the feature vectors for the library entries. For example, a cosine distance metric can be used to compare the input word feature vector against the feature vectors for words in the dictionary which are of the same length. The closest match to the detected input word can then be selected from the dictionary.

**[0068]** In addition to domain specific dictionary models, certain domains can include text fields that follow an expected state transition model or heuristic. For example, in a video of a baseball game, a caption box generally includes a ball-strike count that must follow certain state relationships which can be used to identify and correct character recognition errors. For example, the ball-strike sequence 0-1, 0-2, 1-1 can be recognized as impermissible and flagged as an error for manual editing or corrected using further probabilistic processing.

[0069] To take advantage of such specific transition rules, a transition graph model, such as illustrated in FIG. 9, can be used as a domain specific model. In the transition graph, nodes 900 are assigned a node score and the transitions 905 between nodes assigned a transition score. A node score can be defined as the best correlation value between the input word and the template word in the library. The transition score is defined as the conditional probability:

$$S_{\lambda}(n_{t-1}, n_t) = p(n_t | n_{t-1}) \quad \text{Eq. 2}$$

[0070] where  $n_t$ ,  $n_{t-1}$  are the nodes at the time  $t$  and the time  $t-1$ . The transition conditional probability can be estimated from observations of actual events in the domain, such as actual strike-ball sequences. A small probability is initially assigned to each conditional probability to handle possible misses or false alarms of character detection. A weighting factor  $\lambda$  is then used to combine the node cost and transition cost as an overall cost value at one node:

$$S(n_t) = \lambda S_n(n_t) + (1-\lambda) S_{\lambda}(n_{t-1}, n_t) \quad \text{Eq. 3}$$

Using this model, character recognition can be performed by searching the paths of the transition graph to identify the longest paths. The weighting factor  $\lambda$  can be used to adjust the balance between the contribution from the image based text recognition methods, such as Zernike moments, and domain specific knowledge-based approach. When  $\lambda$  is small, the character recognition process favors the optical character recognition results. As  $\lambda$  increases, the domain based model becomes more significant. The value of  $\lambda$  is chosen experimentally. In the baseball video domain, a value of  $\lambda=0.02$  has been found to provide satisfactory results, with the knowledge based domain model improving the character recognition substantially.

[0071] The extracted caption box text can be used to generate a text based abstract of the video content which indexes the state changes, such as game statistics throughout the content. In addition, the caption box text can also be used as part of an event based video summarization system. An event based video summarization system can use the caption box text to identify events in the video content and then select a portion of the video content based on the detected events. Recognizing that changes in certain word regions in the caption box reflect important events in the video content, a simplified summary of the content can be generated by detecting the events of interest from the caption box text and selecting a video segment of a predetermined length preceding the event. While this approach may lead to satisfactory results, use of a recognized syntactic structure of the event depicted in the video content can result in improved summary generation.

[0072] A summary generation system for baseball video content will now be described to illustrate the use syntactic structure of video content to refine the selection of video segments for inclusion in a summary. Events in a baseball game can be classified into categories, such as score, last pitch (batter change) and base change. A score event is one in which the score for either team changes. A last pitch event is associated with a change in the current batter. This can be the result of a hit where the batter advances to base, a walk where the runner advances to base or an out.

[0073] FIG. 11 is a flow diagram which illustrates a sequence of semantic events that have been consistently observed in baseball video content. The caption change 1110 which is determined from the caption box text extraction and decoding generally represents the end of an event sequence. An event of interest 1120 is preceded by a pitching event

1130. As described in the article "Structure Analysis of Sports Video using Domain Models," by D. Zhong and S. F. Chang, IEEE International Conference on Multimedia and Expo, Aug. 22-25, 2001, Tokyo, Japan, the disclosure of which is hereby incorporated by reference in its entirety, a pitching event is recognizable based on the camera view which generally includes the pitcher and batter in the frame.

[0074] The event of interest 1120 is generally characterized by a change in camera view to one which includes a substantial view of the playing field and a resulting increase in the number of green pixels in the video frame. The event 1120 is generally followed by a view of the player's dugout or the stands to illustrate reaction to the event. This view, referred to as a non-active view 1140, is characterized by a marked decrease in both the motion intensity and the number of green pixels in the frame since the playing field is no longer the predominant feature in the frame. A replay of the event 1150 will generally be presented followed by the change in the caption box reflecting the occurrence of the event. Using this syntactic model, a video segment of the event 1120 can be selected using the transition from the pitching event 1130 as the beginning of the segment and the transition to the non-active view 1140 as the end of the segment. If desired, the pitching event and a portion of the non-active view can also be included in the selected event segment to place the event in context in the resulting summary.

[0075] In order to determine those changes in the caption box which relate to an event of interest, semantic knowledge of the caption box text is beneficial. In this regard, a word region type determination is made for the word regions in the caption box. FIG. 12 illustrates an exemplary caption box for baseball video content and the word types therein. Typical word types include team name 1210, score 1220, inning 1230, out count 1240 and ball-strike count. Domain based knowledge of the video content can be used to determine the word type for the word regions. In general, the word region type determination need be made only once during the video content since the arrangement of the caption box remains constant throughout the content.

[0076] The score word region 1220 is generally represented by numeric characters in the word regions in the caption box which are adjacent to the team name word regions 1210. The team name regions 1210 can be identified using a domain specific library which includes the set of team names and expected abbreviations for such names. The spatial proximity and arrangement of the numeric characters also provides an indication that the numeric characters represent a score type region. For example, as illustrated in FIG. 12, the scores for the two teams generally are vertically aligned with respect to one another.

[0077] The ball-strike region can be identified by the relatively frequent changes in this word region compared to other word regions as well as compliance with an expected state transition which is described above in connection with FIG. 9. Similarly, a transition model recognizing that the out count word region follows the transition 0-1-2-3-0 can be used to identify such a word region.

[0078] In a preferred embodiment, summaries of baseball video content are generated using score and last pitch as events of interest. Changes in score are determined by detecting changes in the score word regions in the caption box text. Last pitch events can be detected using a decision tree model which evaluates changes in ball count and out count to identify the event.

[0079] FIG. 13 illustrates an event decision tree which evaluates changes in caption text to determine if a score or last pitch event have taken place. For each change in the caption box text detected in step 1305, the score type word region is evaluated to determine if a score change event has taken place. If the score has changed in step 1310, then the caption change represents a score event of interest 1315. In step 1320 the ball-strike count region is evaluated to determine if the ball-strike count has been reset. Counts such as 0-0, 0-1 or 1-0 can be used as indicators that the ball-strike count has been reset. The use of 0-1 and 1-0 reduces the possibility of inadvertently overlooking a reset event which may arise due to character recognition errors and the like. If the ball-strike count has been reset, a last pitch event 1330 is detected. If the ball-strike count has not been reset, the out count region is evaluated to determine if the number of outs has changed (step 1340). If so, then a last pitch event is detected in step 1350, otherwise the change in caption text does not represent a last pitch 1355 or score event and is not considered an event of interest. For each caption change representing a score event or a last pitch event, a video segment of the event 1120 can be extracted as described in connection with FIG. 11.

[0080] FIG. 14 illustrates an exemplary screen view of a computer-based video summary browser in accordance with the present invention. The computer-based summary browser is generally implemented in software installed on a conventional computer system, such as an IBM compatible personal computer or Apple Macintosh computer system. In such a system, the video summary content is generally stored in computer readable media, such as a compact disk or magnetic disk storage. A graphical user interface is provided which uses a digital pointer, such as a mouse, to move a selector icon on a display unit in order to control the summary browser.

[0081] The exemplary display of FIG. 14 provides four viewing panels which include a text window 1410 listing scoring event highlights along with a video window 1420 on which a selected scoring highlight video segment can be displayed. A second text window 1430 lists last pitch events adjacent to a video window 1440 for viewing a selected last pitch event. The text windows list the scoring and last pitch events of interest identified in FIG. 13 and include an activation icon 1450 for selecting the video segment associated with that event, such as can be selected using the method of FIG. 11. A user can scroll through the entries in the text window and then select and view those video segments which are of interest to the viewer by clicking on icon 1450 in a manner well known with respect to graphical user interfaces (GUI).

[0082] The specific selection and arrangement of the windows 1410, 1420, 1430 and 1440 in FIG. 14 is not critical to the practice of the present invention. For example, the display can be partitioned into a single text window and a single video window rather than the four windows as shown. In addition, the text window and video window can take on various sizes and shapes as desired by the viewer.

[0083] Although the present invention has been described in connection with specific exemplary embodiments, it should be understood that various changes, substitutions and alterations can be made to the disclosed embodiments without departing from the spirit and scope of the invention as set forth in the appended claims.

1. A system for generating an event based summary of one or more fields or frames of video content including caption boxes embedded therein, comprising:

extracting means for extracting caption boxes from at least a portion of the fields or frames of the video content;  
change identifying means, coupled to the extracting means and receiving the caption boxes therefrom, for identifying one or more changes in the content of the extracted caption boxes which are indicative of an event of interest; and

model applying means, coupled to the extracting means and change identifying means and receiving the caption boxes and the one or more changes therefrom, for applying a semantic model to select a portion of the video content preceding each of the one or more changes.

2. The system of claim 1, wherein the video content comprises video of a baseball game, and wherein the semantic model comprises a model which identifies the portion of the video content of the event of interest as residing between a pitching event and non-active views.

3. The system of claim 2, wherein the events of interest comprise score events and last pitch events.

4. The system of claim 1, wherein the extracting means further comprises:

location means for determining at least one expected location of a caption box in one or more fields or frames of the video content;

determining means, coupled to the location means and receiving the at least one expected location therefrom, for determining at least one caption box mask within the expected location;

frame identifying means, coupled to the determining means and location means and receiving the at least one caption box mask and the at least one expected location therefrom, for identifying one or more fields or frames in the video content as caption fields or frames if the current field or frame exhibits substantial correlation to the at least one caption box mask within the expected caption box;

word region identifying means, coupled to the frame identifying means and receiving the identified caption fields or frames therefrom, for at least a portion of the caption fields or frames, identifying word regions within the confines of the expected location; and

text character means, coupled to the word region identifying means and receiving the word regions therefrom, for each word region, identifying text characters within the region; and processing the identified text characters.

5. The system of claim 4, further comprising comparing means, coupled to the text character means and receiving the text characters therefrom, for comparing the text characters in the word region against a domain specific model to enhance word recognition.

6. The system of claim 4, wherein the location means includes means for evaluating motion features of the video field or frame in the compressed domain, evaluating texture features of the video field or frame in the compressed domain, and identifying regions having low motion features and high texture features as candidate caption box regions.

7. The system of claim 5, wherein the location means includes means for evaluating motion features of the video field or frame in the compressed domain, evaluating texture features of the video field or frame in the compressed domain, and identifying regions having low motion features and high texture features as candidate caption box regions.

8. The system of claim 4, further comprising removal means, coupled to the frame identifying means, the determining means and the location means and receiving the identified caption fields or frames, the at least one caption box mask and the at least one expected location therefrom, for evaluating the identified caption fields or frames, within the caption box location, for changes in content and removing caption fields or frames from word region processing which do not exhibit a change in content.

9. The system of claim 4, further comprising interval means, coupled to the frame identifying means and receiving the identified caption fields or frames therefrom, for selecting a subset of caption fields or frames based on a predetermined time interval and sending that subset to the word region identifying means.

10. The system of claim 4, further comprising number means, coupled to the frame identifying means and receiving the identified caption fields or frames therefrom, for determining a subset of caption fields or frames by selecting caption fields or frames based on a predetermined number of intervening caption fields or frames and sending that subset to the word region identifying means.

11. A non-transitory computer-readable storage medium storing a program for causing a computer to implement a method for generating an event based summary of one or more fields or frames of video content including caption boxes embedded therein, comprising:

- extracting caption boxes from at least a portion of the fields or frames of the video content;
- identifying one or more changes in the content of the extracted caption boxes which are indicative of an event of interest; and
- applying a semantic model to select a portion of the video content preceding each of the one or more changes.

12. The non-transitory computer-readable storage medium of claim 11, wherein the extracting caption boxes further comprises:

- determining at least one expected location of a caption box in one or more fields or frames of the video content;
- determining at least one caption box mask within the expected location;
- identifying one or more fields or frames in the video content as caption fields or frames if the current field or frame exhibits substantial correlation to the at least one caption box mask within the expected caption box location;
- for at least a portion of the caption fields or frames, identifying word regions within the confines of the expected location;
- for each word region, identifying text characters within the region, and processing the identified text characters.

13. The non-transitory computer-readable storage medium of claim 12, wherein the extracting caption boxes includes the step of comparing the text characters in the word region against a domain specific model to enhance word recognition.

14. A system for generating an event based summary of one or more fields or frames of video content including caption boxes embedded therein, comprising:

- a caption box extractor for extracting caption boxes from at least a portion of the fields or frames of the video content;
- a caption change identification processor, coupled to the caption box extractor and receiving the caption boxes therefrom, for identifying one or more changes in the

content of the extracted caption boxes which are indicative of an event of interest; and

- a model processor, coupled to the caption box extractor and caption change identification processor and receiving the caption boxes and the one or more changes therefrom, for applying a semantic model to select a portion of the video content preceding each of the one or more changes.

15. The system of claim 14, wherein the video content comprises video of a baseball game, and wherein the semantic model comprises a model which identifies the portion of the video content of the event of interest as residing between a pitching event and non-active views.

16. The system of claim 14, wherein the event of interest comprises one or more score events and last pitch events.

17. The system of claim 14, wherein the caption box extractor further comprises:

- a location processor for determining at least one expected location of a caption box in one or more fields or frames of the video content;
- a caption mask determining processor, coupled to the location processor and receiving the at least one expected location therefrom, for determining at least one caption box mask within the expected location;
- a frame identification processor, coupled to the caption mask determining processor and location processor and receiving the at least one caption box mask and the at least one expected location therefrom, for identifying one or more fields or frames in the video content as caption fields or frames if the current field or frame exhibits substantial correlation to the at least one caption box mask within the expected caption box;
- a word region identification processor, coupled to the frame identification processor and receiving the identified caption fields or frames therefrom, for at least a portion of the caption fields or frames, identifying word regions within the confines of the expected location; and
- a text character identification processor, coupled to the word region identification processor and receiving the word regions therefrom, for each word region, identifying text characters within the region; and processing the identified text characters.

18. The system of claim 17, further comprising a text character comparison processor, coupled to the text character identification processor and receiving the text characters therefrom, for comparing the text characters in the word region against a domain specific model to enhance word recognition.

19. The system of claim 17, wherein the location processor includes a processor for evaluating motion features of the video field or frame in the compressed domain, evaluating texture features of the video field or frame in the compressed domain, and identifying regions having low motion features and high texture features as candidate caption box regions.

20. The system of claim 18, wherein the location processor includes a processor for evaluating motion features of the video field or frame in the compressed domain, evaluating texture features of the video field or frame in the compressed domain, and identifying regions having low motion features and high texture features as candidate caption box regions.

21. The system of claim 17, further comprising a caption removal processor, coupled to the frame identification processor, the caption mask determining processor and the location processor and receiving the identified caption fields or

frames, the at least one caption box mask and the at least one expected location therefrom, for evaluating the identified caption fields or frames, within the caption box location, for changes in content and removing caption fields or frames from word region processing which do not exhibit a change in content.

**22.** The system of claim **17**, further comprising an interval-based caption selection processor, coupled to the frame identification processor and receiving the identified caption fields or frames therefrom, for selecting a subset of caption fields or frames based on a predetermined time interval and sending that subset to the word region identification processor.

**23.** The system of claim **17**, further comprising a number-based caption selection processor, coupled to the frame identification processor and receiving the identified caption fields or frames therefrom, for determining a subset of caption fields or frames by selecting caption fields or frames based on a predetermined number of intervening caption fields or frames and sending that subset to the word region identification processor.

\* \* \* \* \*