



(51) International Patent Classification:
G06K 9/00 (2006.01)

(21) International Application Number:

PCT/US2017/032117

(22) International Filing Date:

11 May 2017 (11.05.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

15/153,236

12 May 2016 (12.05.2016)

US

(71) Applicant: **ARRIS ENTERPRISES LLC** [US/US]; 3871
Lakefield Drive, Suwanee, Georgia 30024 (US).

(72) Inventors: **EMEOTT, Stephen P.**; 2616 Moray Lane,
Cedar Bank, Texas 78613 (US). **LI, Renxiang**; 290 Stone
Avenue, Lake Zurich, Illinois 60047 (US).

(74) Agent: **WIENER, Stewart M.** et al.; 3871 Lakefield Drive,
Suwanee, Georgia 30024 (US).

(81) Designated States (*unless otherwise indicated, for every
kind of national protection available*): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR,
KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG,
MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM,
PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC,
SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR,
TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(54) Title: DETECTING SENTINEL FRAMES IN VIDEO DELIVERY USING A PATTERN ANALYSIS




FIG. 3

(57) Abstract: Particular embodiments extract features for frames from a video. Then, a sequence of frames is identified based on a pattern analysis of frames based on the features. This pattern may be used to select sentinel features from the sequence of frames. For example, a pattern may include a first sentinel sequence followed by a transitional frame and then a second sentinel sequence. The transitional frame may be a black frame that is used to identify back-to-back sentinel sequences. Then, a sentinel frame that demarks a transition from a first content type to a second content type is identified. For example, the frames on either side of the black frame may be very similar and be identified as sentinel frames that mark a transition from program content to advertisement content. The above process may allow automated detection of sentinel frames that can run without user supervision.



(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

DETECTING SENTINEL FRAMES IN VIDEO DELIVERY USING A PATTERN ANALYSIS

BACKGROUND

[0001] Sentinel sequences may identify transitions between a first type of content and a second type of content. These sequences may also be referred to as “bumper” sequences. Sentinel sequences may be short clips or animations that typically last a short period of time, such as 2 - 5 seconds, and can be replayed several times in the same video. Also, similar sentinel sequences may be displayed in other videos over a time period, such as over the course of a day or longer.

[0002] The sentinel sequences may be identified using one or more key frames that are manually extracted from the recorded content. For example, a user may manually view the content and identify the sentinel sequences. Then, the user needs to manually extract a frame from the sentinel sequence. This may help an ad detection system to replace the ads. However, this is a highly manual process, and it is hard to train a user to extract the correct sentinel frame.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] FIG. 1 depicts a simplified system for identifying sentinel frames according to one embodiment.

[0004] FIG. 2 depicts an example of a pattern that could be used to determine sentinel frames according to one embodiment.

[0005] FIG. 3 depicts a more detailed example of a sentinel frame processor according to one embodiment.

[0006] FIG. 4 depicts a simplified flowchart of a method for generating sentinel frames according to one embodiment.

[0007] FIG. 5 depicts a simplified flowchart of an example sequential clustering process that may be used according to one embodiment.

[0008] FIG. 6 depicts an example of a non-sequential clustering process according to one embodiment.

[0009] FIG. 7 depicts a flowchart of the overall process for using a sentinel frame data set according to one embodiment.

[0010] FIG. 8 illustrates an example of a special purpose computer system configured with the sentinel frame processor according to one embodiment.

DETAILED DESCRIPTION

[0011] Described herein are techniques for a content detection system. In the following description, for purposes of explanation, numerous examples and specific details are set forth in order to provide a thorough understanding of particular embodiments. Particular embodiments as defined by the claims may include some or all of the features in these examples alone or in combination with other features described below, and may further include modifications and equivalents of the features and concepts described herein.

[0012] Particular embodiments extract features for frames from a video. Then, a sequence of frames is identified based on a pattern analysis of frames based on the features. This pattern may be used to select sentinel features from the sequence of frames. For example, a pattern may include a first sentinel sequence followed by a transitional frame and then a second sentinel sequence. The transitional frame may be a black frame that is used to identify back-to-back sentinel sequences. Then, a sentinel frame that demarks a transition from a first content type to a second content type is identified. For example, the frames on either side of the black frame may be very similar

and be identified as sentinel frames that mark a transition from program content to advertisement content. The above process may allow automated detection of sentinel frames that can run without user supervision.

Overview

[0013] FIG. 1 depicts a simplified system 100 for identifying sentinel frames according to one embodiment. System 100 includes a video processor 102, a video service 104, a video delivery system 106, and clients 108. Video processor 102 may receive a video and/or multiple videos from video content providers. Video processor 102 may be part of a video service provider that may deliver the videos to clients 108. For example, clients 108 may include customer premise equipment (CPE), such as set top boxes, mobile devices, or other computing devices, that receive the videos, such as through a broadcast or on-demand in a cable network.

[0014] Video processor 102 includes a sentinel frame processor 110 that can analyze the video to determine sentinel frames. The sentinel frames are identified without prior knowledge of the content of the sentinel frames. The sentinel frames signal a transition or boundary between a first type of content and a second type of content. For example, the sentinel frames may be key frames or images that may be repeatedly used to identify ad boundaries. This may be where program content transitions to advertisement content, different types of ads are shown, or a segment in a show transitions to a new segment. As will be discussed in more detail below, a pattern for frames of the video may be used to identify sentinel frames.

[0015] Video processor 102 generates visual features for the sentinel frames. Although visual features may be used for discussion purposes, other features may be used, such as audio, textual, and/or different combinations of audio, visual, and textual. This allows video service 104 to search videos to identify where sentinel frames appear in the videos.

Video service 104 may then perform services on the video program using the locations of the sentinel frames.

[0016] Once the position of sentinel frames in the video are identified, video service 104 may perform a service, such as replacing the video segments with different video content. For example, video service 104 may replace advertisements from one locality for advertisements that are more pertinent to a second locality. Although video replacement is described, other services may be performed, such as translation and video editing.

[0017] Video service 104 then passes the video with the replaced content to video delivery system 106. Video delivery system 106 may then be used to deliver the video program with the replaced video content to clients 108. In one embodiment, video delivery system 106 may have the video content replaced in substantially real time. For example, the sentinel frames may be detected and video content replaced by video service 104 with only a one or two second delay in delivering video content from a broadcast (e.g., a live or scheduled broadcast). In other embodiments, video delivery system 106 may have the video replaced in the video content in the background or before the initiation of delivery of the video program. In other embodiments, the video delivery system 106 may extract a video program from a live video stream based upon sentinel frames located at the beginning and the end of the program.

Example pattern to identify sentinel frames

[0018] As mentioned above, sentinel frame processor 110 may use patterns to detect sentinel frames. FIG. 2 depicts an example of a pattern that could be used to determine sentinel frames according to one embodiment. In one embodiment, the pattern is on the frame level. For example, the pattern is based on frame-level sequences between the frames and not based on content within a single frame or a within portion of a single frame. For example, the pattern may be a sequence of sentinel frames followed by a

transition frame and then a sequence of sentinel frames again. The sentinel frame is not detected based on a logo or something else found within a single sentinel frame.

[0019] A video includes a series of frames 202-1 - 202-5 labeled in sequence from frames 21562-21566 of the video. There are other frames that appear before and after this sequence in the video (not shown). At 204-1, a first sentinel frame sequence is shown and at 204-2, a second sentinel frame sequence is shown. The sentinel frame sequence shown at 204-1 may be the end of a first sentinel frame sequence and the sentinel frame sequence shown at 204-2 may be the start of a second sentinel frame sequence. The first and second sentinel sequences may be the same sequence or different sequences. However, the first and second sentinel sequences include at least one frame that is similar.

[0020] A transition frame 202-3, such as a black frame, may be different from the sentinel frame sequences. This transition frame may be used to signal a transition from the first sentinel frame sequence to the second sentinel frame sequence. Examples of sentinel transitional frames may include black frames, animated sequences, other solid color frames, or other information that may be inserted between multiple sentinel frame sequences. These transitional frames may be associated between two sentinel sequences. In this case, the pattern may be back-to-back sentinel frame sequences in a video that includes a transition frame between the back-to-back sentinel frame sequences. The back-to-back sentinel frame sequences may be during a transition from sponsor ads to product ads, or from program content to advertising content. The one or more transition frames, such as black frames, may be in between the sentinel frame sequence and may be identified. For example, black frames may be easily identified based on visual features of the frame such as, a color layout descriptor (CLD) or edge histogram (EHD) may be used to detect black frames.

[0021] When back-to-back sentinel frame sequences are used, the frames on either side of the transition frame may be very similar. For example, frame 202-2 is similar to frame 202-4. This pattern may be used to identify sentinel frames. Although this pattern is

discussed, it will be understood that other patterns may also be used. For example, a one-sided sentinel frame sequence may occur when a video program transitions from program to advertisements and back again. The one-sided pattern occurs when video content on a first side of the black frame is a sentinel frame sequence and the video content on the second side is not a sentinel frame sequence. Rather, the video on the second side may be an advertisement. This pattern is detectable because the sentinel frame sequence forms a pattern that will be repeated at many program breaks whereas the non-sentinel frame sequence on the other side of the black frame will be repeated far less often. However, a one-sided sentinel frame sequence may require a larger number of observations to detect the pattern and sentinel frame than the back-to-back pattern due to the statistical uncertainties about whether a black frame adjacent to two dissimilar sequences represents a program break including a sentinel frame sequence or an unmarked scene transition. A program break sentinel frame sequence may optionally be dependent upon program metadata, such as the program title, such as when the one-sided sentinel frame pattern is unique to the program. The metadata may be leveraged in detecting the pattern.

[0022] Another example of a sentinel pattern is a promotional segment sentinel frame sequence. A pattern for a promotional segment sentinel frame sequence may occur at the beginning or end of a promotional video inserted by a broadcaster to promote an upcoming event or broadcast. This sentinel may be a one sided sentinel frame sequence, and may be adjacent to a black frame. The promotional segment may also be identified by closed captions associated with the promotion. The captions help identify the promotional segment, but a transition event, such as black frames or a hard cut, is used to identify the boundaries of the promotional segment sentinel frame series.

[0023] Program start and program end sentinel frame sequences are another pattern that may be identified, and are similar to promotional segment sentinel frame series. Promotional segment sentinel frame series are repeated frequently in a pattern, sometime at multiple instances throughout a program. Program start and end sentinel frame

sequences are repeated much less frequently, but program metadata information and electronic program guide information may be used to correlate likely program start and program end sentinel frame sequences with similar sequences from other episodes of the program. Not every program will have repeating program start and program end sentinel frame sequences but those that do may be detected using the pattern. The process of detecting these patterns will now be discussed in more detail.

Sentinel Frame Processor

[0024] FIG. 3 depicts a more detailed example of sentinel frame processor 110 according to one embodiment. A feature extraction service 302 receives the video and can extract features from the frames of the video. Although a single video is described, multiple videos may be processed. The features that could be extracted may be a color layout descriptor (CLD) or other visual features, such as edges that form an edge histogram. Although these visual features are described, other features may also be extracted. The color layout descriptor describes a distribution of color in a frame and the edge histogram summarizes edges found in the frame. The values for these features may be determined based on decoding a frame of the video. Accordingly, feature extraction service 302 may include logic for a decoder (not shown) or receive these features from a decoder. In another embodiment, the feature extraction service 302 uses optical character recognition to identify on-screen text marking program transitions. For example, the presence of on-screen text correlated with the program may mark the beginning or end of a program, or the beginning or end of a program break. In some embodiments, the on-screen text is known a priori, while in other embodiments the text is unpredictable while the underlying video background is predictable. In another embodiment, the feature extraction service may analyze the video and detect transitions where on-screen logos or ticker information appears and disappears. When combined with video transitions, such as fades or hard cuts, these transitions may be indicative of the start or end of a sentinel frame sequence and used as a pattern.

[0025] Pattern analysis service 304 may receive user input identifying patterns to look for. In other embodiments, pattern analysis service 304 may dynamically analyze a video to identify patterns that could be used. For example, a system administrator might be presented with a list of checkboxes identifying features known by the system and likely to occur near sentinel frame sequences. In one embodiment, the list might include black frames, a visual channel logo, and electronic guide information associated with one or more program. The system administrator may select one or more of the features to use as a trigger for sentinel frame identification. In another embodiment, patterns common to several days or months of programming may be automatically set to trigger sentinel frame identification when the number of sentinel frame detected over some number of hours or days goes below a threshold. In this embodiment, the relative absence of detected sentinels is indicative of changes to the underlying sentinel frame sequences.

[0026] In one embodiment, pattern analysis service 304 first uses the visual features to identify the location of black frames. Then, pattern analysis service 304 organizes sequential black frames into groups. That is, the black frame transition may last for a couple of frames in the video. For each group of black frames that have a number of frames that is under a threshold, such as groups of black frames of the length 1 - 3 frames, pattern analysis service 304 compares the visual features of non-black frames on either side of the black frame group to identify similar frame pairs. For example, this identifies two similar frames on both sides of the black frame. These frames may be right next to the black frame group or within a number of frames (e.g., within a threshold such as 5 frames).

[0027] Pattern analysis service 304 determines that frames are similar based on a threshold. For example, two frames may be similar if they have feature values that are similar within a threshold. When a similar frame-black frame-similar frame sequence is found, pattern analysis service 304 notes the location of the pattern. If pattern analysis service 304 does not find two frames that are similar on both sides of the black frame, pattern analysis service 304 does not consider this sequence a pattern.

[0028] After reviewing all the groups of black frames, pattern analysis service 304 may record the location of the patterns found. For example, the location of sentinel frame-black frame-sentinel frame patterns is recorded. Each pattern may be the sentinel sequence, black frame, sentinel sequence pattern.

[0029] A sentinel frame identification service 306 can then identify sentinel frames from the identified patterns. For example, sentinel frame identification service 306 may extract visual features for two or more sentinel frames in each pattern to create a candidate sentinel frame data set.

[0030] Sentinel frame identification service 306 then clusters sentinel frames in the candidate sentinel data set into one or more groups based on the visual features. For example, sentinel frames that may be similar may be clustered into a group. Then, sentinel frame identification service 306 returns a representative sentinel frame for each group. That is, the visual features for the frames in each group may be used to generate a sentinel frame.

[0031] Sentinel frame identification service 306 may generate the sentinel frame in different ways, such as by selecting a sentinel frame that is closest to an average the values for the visual features in the group. In other embodiments, an average of the visual features in the group may be used as the sentinel frame. In one example, the color layout descriptor (CLD) values may be generated for a sentinel frame based on the color layout descriptor values found in the group. Sentinel frame identification service 306 then outputs the locations of the sentinel frames in the video and the values for the sentinel frames. An example of the color layout descriptor for a sentinel frame may be [58, 19, 18, 16, 17, 15, 31, 17, 16, 32, 16, 16]. These numbers of the color layout descriptor capture the spatial distribution of color in an image. For example, the color layout descriptor values capture the spatial layout of the representative colors on a grid superimposed on a region or image. The representation is based on coefficients of the discrete cosine transform (DCT) and is a way of encoding RGB (Red, Green, Blue) information of an image, where the first 6 integer values represent the luma component of

the CLD, the next 3 integer values represent the blue difference chroma component of the CLD and the final 3 integer values represent the red difference chroma component of the CLD. The above values represent the coefficients for the color layout descriptor for one exemplary sentinel frame.

[0032] FIG. 4 depicts a simplified flowchart 400 of a method for generating sentinel frames according to one embodiment. At 402, sentinel frame processor 110 identifies a pattern. The pattern may be identified based on receiving user input regarding the pattern, may be identified dynamically while processing the video or was identified from prior videos. At 404, sentinel frame processor 110 identifies patterns within a video. For example, after finding black frames, sentinel frame processor 110 determines if similar frames are located on each side of the black frames.

[0033] At 406, sentinel frame processor 110 extracts visual features for the sentinel frames in the patterns. For example, in the pattern described above, sentinel frame processor 110 identifies visual features for sentinel frames on both sides of a black frame. The visual features may include color layout descriptor values.

[0034] At 408, sentinel frame processor 110 creates a candidate sentinel data set. The candidate sentinel data set includes all of the candidate sentinel frames.

[0035] At 410, sentinel frame processor 110 clusters frames in the candidate sentinel data set into groups based on the visual features. This creates groups of possible sentinel frames that are similar within a threshold. At 412, sentinel frame processor 110 then selects sentinel frames for each group.

Clustering Processes

[0036] In steps 410 and 412, clustering was discussed. In one embodiment, sentinel frame processor 110 may use different clustering processes to generate sentinel frames. For example, sentinel frame processor 110 can use sequential and non-sequential

processes. Although these clustering processes are used, it will be understood that other clustering processes may be used.

[0037] FIG. 5 depicts a simplified flowchart 500 of an example sequential clustering process that may be used according to one embodiment. In this example, a sequential process is used as sentinel frames are generated sequentially, and not in parallel.

[0038] At 502, sentinel frame processor 110 selects a random sentinel frame from the candidate sentinel data set. For example, the first sentinel frame in the data set may be selected, but other sentinel frames may also be selected. For the remaining sentinel frames in the candidate data set, at 504, sentinel frame processor 110 creates a first “in group” for sentinels that have visual features within a threshold close to the selected sentinel. For example, a distance between the visual features, such as a CLD vector, can be defined as a maximum distance between elements in the selected sentinel frame CLD vector and the candidate frame CLD vectors.

[0039] At 506, for the remaining sentinels in the candidate data set, sentinel frame processor 110 creates an “out group” for the candidate sentinel frames that have values that are above the threshold to the selected sentinel. For example, candidate sentinel frames that have visual features with a large distance between elements in the candidate sentinel frame and the selected sentinel frame are placed in the out group. Accordingly, the out group is for sentinels that are not similar to the selected sentinel frame and the in group is for sentinel frames that are similar to the selected sentinel.

[0040] At 508, sentinel frame processor 110 averages the sentinel frames in the in group to generate an average sentinel frame. This may average the visual feature values for all of the in group sentinel frames.

[0041] At 510, sentinel frame processor 110 creates a second in group for sentinel frames close to the average sentinel frame. Then, at 512, sentinel frame processor 110 creates an out group for sentinel frames that are not close to the average sentinel frame.

[0042] At 514, sentinel frame processor 110 determines if a sentinel frame from the second in group exists. Sentinel frame processor 110 may output no sentinel frames if the number of sentinel frames in the second in group is less than a threshold. The candidate sentinel data set often includes a number of non-sentinel frames, as while many sentinel frames are located adjacent to transition frames not every transition frame is located adjacent to a sentinel frame. Cases where the number of sentinel frames in the second in group is below a threshold represent instances where the random sentinel frame selected in 502 is not a sentinel frame. In these cases, at 516, sentinel frame processor 110 does not output a representative sentinel frame. If a sentinel frame exists, at 518, sentinel frame processor 110 selects a sentinel frame from the second in group that is closest to the average in group sentinel and outputs that as a representative sentinel frame. This representative sentinel frame can then be used to detect sentinel frames in videos. In other embodiments, the average sentinel frame may be selected as the representative sentinel frame. However, because the average is being used, there may be slight differences between the average sentinel frame and sentinel frames found in the videos.

[0043] At 520, sentinel frame processor 110 determines if another iteration should be performed. If so, the above process is then repeated beginning at 502 with the sentinel frames in the out group being the sentinel frames in the candidate sentinel data set. If not, at the end, sentinel frame processor 110 has processed all the candidate sentinel frames and output representative sentinel frames that were found in the video.

[0044] FIG. 6 depicts an example of a non-sequential clustering process according to one embodiment. Non-sequential may mean that the process is not performed by analyzing frames sequentially, but rather in parallel.

[0045] At 602, sentinel frame processor 110 uses a process, such as a K-means or similar algorithm, to cluster frames in the candidate sentinel frame data set to create candidate clusters. For example, this may initially start with N clusters, where N is the expected number of different sentinel frame sequences. In one embodiment, different

sentinel frame sequences are expected for different programming categories, such as prime time programming, children's programming, sports programming and news programming. In this embodiment $N=4$. Another embodiment may initially start with N clusters, where N is the number of video programs included in the analysis. Since the number of sentinel frame sequences is typically smaller than the number of video programs included in the analysis, the number of candidate clusters are pared down as described below.

[0046] At 604, sentinel frame processor 110 removes clusters that include a number of elements that is below a threshold to produce a set of valid clusters. For example, the clusters removed may be clusters with a low number of elements, such as with no elements or with 1 element. A cluster with no elements may be a video program that did not include any unique candidate sentinel frames.

[0047] At 606, for each of the remaining M clusters (e.g., $M \leq N$), sentinel frame processor 110 measures the maximum element distance and the distance between cluster centers for all clusters. The cluster center may be a value that is in the center of the values of the cluster. This measures how similar clusters are using the distance from the center of the clusters and the maximum distance between elements in the clusters. Although the center is discussed, other points in the clusters may be used.

[0048] At 608, sentinel frame processor 110 merges nearby clusters. For example, clusters having center distances that are less than a threshold T_d may be merged. The threshold T_d may be equal to $(1 + \alpha) * (\text{maximum in-cluster element distance for all clusters})$, where α is typically a positive real number less than one. This means that if the average distance between elements within valid clusters is larger than the distance between the centers of two candidate clusters, then the two candidate clusters may be merged.

[0049] At 610, sentinel frame processor 110 selects a sentinel frame from each remaining cluster. This is because all similar sentinel frames have been merged. In particular, when the average distance between elements in a cluster is significantly

smaller than the distance between cluster centers and when the size of each cluster exceeds a threshold, then the number of clusters remaining is unlikely to exceed the number of unique sentinel frames. Thus, each of the remaining clusters is a valid cluster and the sentinel frame processor 110 may proceed to selecting a sentinel frame. . For example, sentinel frame processor 110 selects one feature that is closest to the cluster center to represent the resulting sentinel frame. For example, the feature value for the center of the cluster may be selected as the resulting sentinel frame. Then, at 612, sentinel frame processor 110 outputs the sentinel frames for the clusters.

Sentinel frame update

[0050] The sentinel frames that are detected may be used to detect sentinel frames in other video programs. Particular embodiments may maintain a database that can include an updated list of sentinel frames to use. FIG. 7 depicts a flowchart 700 of the overall process for using a sentinel frame data set according to one embodiment. At 702, sentinel frame processor 110 detects sentinel frames in a first video. Any of the processes described above may be used.

[0051] At 704, sentinel frame processor 110 generates a sentinel database for all of the video programs. This includes sentinel frames that are detected from all of the video programs.

[0052] In one embodiment, different video programs on a same channel may use different sentinel frames. In other embodiments, the same sentinel frames may be used in different video programs. That is, different shows may use different sentinel frames. However, within episodes of the same show, the same sentinel frames may be used. Also, different sentinel frames may be used on different days or times of the year or for different genres of video programming. For example, prime time programming may employ a first set of sentinel frames while children's programming may employ a second set of sentinel frames.

[0053] At 706, sentinel frame processor 110 determines if any changes are needed to the sentinel frame database based on a new video sentinel frame analysis. For example, sentinel frames may be added or deleted from the current sentinel frame database at any time. If the new sentinel frame sequences start appearing less often or are not detected at all, a sentinel frame update may be needed. Also, if a new sentinel frame sequence starts appearing, then a sentinel frame update may be needed. Thus, using particular embodiments, if a representative sentinel frame is close to a sentinel frame already in the current sentinel frame database, then an update may not be necessary. However, if a representative sentinel frame is unique and not found in the sentinel frame database, then an update of the current sentinel frame database would add the new sentinel frame and could improve ad detection and replacement performance.

[0054] Accordingly, particular embodiments are able to identify sentinel frames based on pattern analysis. The blind detection of sentinel frames allows the system to detect sentinel frames even in the event that different programs use different sentinel frames or in the presence of seasonal changes to the sentinel frames. In this case, even though the sentinel frames change, the overall pattern may be the same and may still be detectable. For example, when a sentinel frame is manually uploaded, particular embodiments allow an operator to validate the presence of the sentinel sequences containing the uploaded sentinel frame using a high-quality metric. That is, sentinel frame processor 110 can detect sentinel frames in a video that should match the uploaded sentinel frame.

Example computer system

[0055] FIG. 8 illustrates an example of a special purpose computer system 800 configured with sentinel frame processor 110 according to one embodiment. In one embodiment, computer system 800-1 describes system 100 including video processor 102. Also, computer system 800-2 describes clients 108. Only one instance of computer system 800 will be described for discussion purposes, but it will be recognized that computer system 800 may be implemented for other entities described above.

[0056] Computer system 800 includes a bus 802, network interface 804, a computer processor 806, a memory 808, a storage device 810, and a display 812.

[0057] Bus 802 may be a communication mechanism for communicating information. Computer processor 806 may execute computer programs stored in memory 808 or storage device 808. Any suitable programming language can be used to implement the routines of particular embodiments including C, C++, Java, assembly language, etc. Different programming techniques can be employed such as procedural or object oriented. The routines can execute on a single computer system 800 or multiple computer systems 800. Further, multiple computer processors 806 may be used.

[0058] Memory 808 may store instructions, such as source code or binary code, for performing the techniques described above. Memory 808 may also be used for storing variables or other intermediate information during execution of instructions to be executed by processor 806. Examples of memory 808 include random access memory (RAM), read only memory (ROM), or both.

[0059] Storage device 810 may also store instructions, such as source code or binary code, for performing the techniques described above. Storage device 810 may additionally store data used and manipulated by computer processor 806. For example, storage device 810 may be a database that is accessed by computer system 800. Other examples of storage device 810 include random access memory (RAM), read only memory (ROM), a hard drive, a magnetic disk, an optical disk, a CD-ROM, a DVD, a flash memory, a USB memory card, or any other medium from which a computer can read.

[0060] Memory 808 or storage device 810 may be an example of a non-transitory computer-readable storage medium for use by or in connection with computer system 800. The non-transitory computer-readable storage medium contains instructions for controlling a computer system 800 to be configured to perform functions described by particular embodiments. The instructions, when executed by one or more computer

processors 806, may be configured to perform that which is described in particular embodiments.

[0061] Computer system 800 includes a display 812 for displaying information to a computer user. Display 812 may display a user interface used by a user to interact with computer system 800.

[0062] Computer system 800 also includes a network interface 804 to provide data communication connection over a network, such as a local area network (LAN) or wide area network (WAN). Wireless networks may also be used. In any such implementation, network interface 804 sends and receives electrical, electromagnetic, or optical signals that carry digital data streams representing various types of information.

[0063] Computer system 800 can send and receive information through network interface 804 across a network 814, which may be an Intranet or the Internet. Computer system 800 may interact with other computer systems 800 through network 814. In some examples, client-server communications occur through network 814. Also, implementations of particular embodiments may be distributed across computer systems 800 through network 814.

[0064] Particular embodiments may be implemented in a non-transitory computer-readable storage medium for use by or in connection with the instruction execution system, apparatus, system, or machine. The computer-readable storage medium contains instructions for controlling a computer system to perform a method described by particular embodiments. The computer system may include one or more computing devices. The instructions, when executed by one or more computer processors, may be configured to perform that which is described in particular embodiments.

[0065] As used in the description herein and throughout the claims that follow, “a”, “an”, and “the” includes plural references unless the context clearly dictates otherwise. Also, as used in the description herein and throughout the claims that follow, the meaning of “in” includes “in” and “on” unless the context clearly dictates otherwise.

[0066] The above description illustrates various embodiments along with examples of how aspects of particular embodiments may be implemented. The above examples and embodiments should not be deemed to be the only embodiments, and are presented to illustrate the flexibility and advantages of particular embodiments as defined by the following claims. Based on the above disclosure and the following claims, other arrangements, embodiments, implementations and equivalents may be employed without departing from the scope hereof as defined by the claims.

CLAIMS

What is claimed is:

1. A method comprising:

extracting, by a computing device, frame features for a plurality of frames from a video;

identifying, by the computing device, a pattern from a sequence of frames of the plurality of frames, the pattern being identified based on a pattern analysis using the frame features for frames in the sequence of frames;

selecting, by the computing device, sentinel features from the frame features for frames in the sequence of frames based on the pattern; and

outputting, by the computing device, a sentinel frame that identifies a transition in the video from a first content type to a second content type, wherein the sentinel frame includes the sentinel features.

2. The method of claim 1, wherein identifying the pattern comprises:

selecting a transitional frame in the sequence of frames, wherein the transitional frame is different from the sentinel frame; and

using at least one frame in the sequence of frames that is before or after the transitional frame to select the sentinel features.

3. The method of claim 2, wherein selecting the sentinel features comprises using a frame in the sequence of frames on both sides of the transitional frame in the sequence of frames to select the sentinel features.

4. The method of claim 2, wherein identifying the pattern comprises:

comparing a first frame on a first side of the transitional frame to a second frame on a second side of the transitional frame to determine if the first frame and the second frame are similar within a threshold; and

when the first frame and the second frame are similar within the threshold, identifying the sequence of frames as including the pattern.

5. The method of claim 2, further comprising:

grouping sequential transitional frames in the sequence of frames into a group, wherein the at least one frame in the sequence of frames is near the group within a threshold.

6. The method of claim 2, wherein the transitional frame is a solid color frame.

7. The method of claim 1, wherein the pattern comprises a first sequence of frames and a transitional frame.

8. The method of claim 7, further comprising:

locating a second sequence of frames in addition to the first sequence of frames and the transitional frame, wherein:

a first frame in the first sequence of frames matches a second frame in the second sequence of frames within a threshold, and

the sentinel features are selected from the first frame or the second frame.

9. The method of claim 1, wherein the sentinel features comprise features determined from decoding the frames of the video.

10. The method of claim 1, further comprising:
receiving an input identifying the pattern.

11. The method of claim 1, wherein the set of sentinel frames are not previously known before selecting the sentinel features.

12. A method comprising:
extracting, by a computing device, frame features for a plurality of frames from a video;
identifying, by the computing device, locations of first frames in the video using the frame features, the first frames including a first set of frame features;
using, by the computing device, frame features for second frames that are within a threshold distance to the first frames to generate a set of candidate frames;
clustering, by the computing device, the set of candidate frames into one or more groups based on the frame features for the set of candidate frames;

selecting, by the computing device, sentinel features for each of the one or more groups; and

outputting, by the computing device, a set of sentinel frames for each of the one or more groups using the sentinel features, the set of sentinel frames identifying a transition in the video from a first content type to a second content type.

13. The method of claim 12, wherein using the frame features for the second frames comprises:

for each first frame, comparing a second frame on at least one side of the first frame to identify a pattern of the second frame and the first frame.

14. The method of claim 12, wherein using the frame features for the second frames comprises:

for each first frame, comparing a second frame on each side of the first frame to generate a difference between the second frame on each side of the first frames; and

adding the second frame on each side of the first frame to the set of candidate frames if the second frame on each side of the first frame are similar within a threshold.

15. The method of claim 12, wherein clustering comprises using a sequential clustering process.

16. The method of claim 15, wherein the sequential clustering process comprises:

selecting a first candidate frame in the set of candidate frames;

creating a first group of candidate frames that include frame features that are similar to the first candidate frame within a similarity threshold;

generating an average sentinel frame based on an average of the frame features of the first group of candidate frames;

selecting a candidate frame in the first group of candidate frames that is closest to the average sentinel frame; and

including the candidate frame in the set of sentinel frames.

17. The method of claim 12, wherein clustering comprises using a non-sequential clustering process.

18. The method of claim 17, wherein the non-sequential clustering process comprises:

clustering candidate frames in the set of candidate frames into a first set of clusters;

merging one or more pluralities of clusters, from the first set of clusters, that have frame features within a similarity threshold, to generate a second set of clusters; and

selecting a sentinel frame from each of the second set of clusters using a frame feature within each of the second set of clusters.

19. The method of claim 12, wherein the set of sentinel frames are not previously known before selecting the sentinel features.

20. An apparatus comprising:

one or more computer processors; and

a non-transitory computer-readable storage medium comprising instructions that, when executed, control the one or more computer processors to be configured for:

extracting frame features for a plurality of frames from a video;

identifying a pattern from a sequence of frames of the plurality of frames, the pattern being identified based on a pattern analysis using the frame features for frames in the sequence of frames;

selecting sentinel features from the frame features for frames in the sequence of frames based on the pattern; and

outputting a sentinel frame that identifies a transition in the video from a first content type to a second content type, wherein the sentinel frame includes the sentinel features.

100




FIG. 1




FIG. 2

3/8




FIG. 3

4/8




FIG. 4

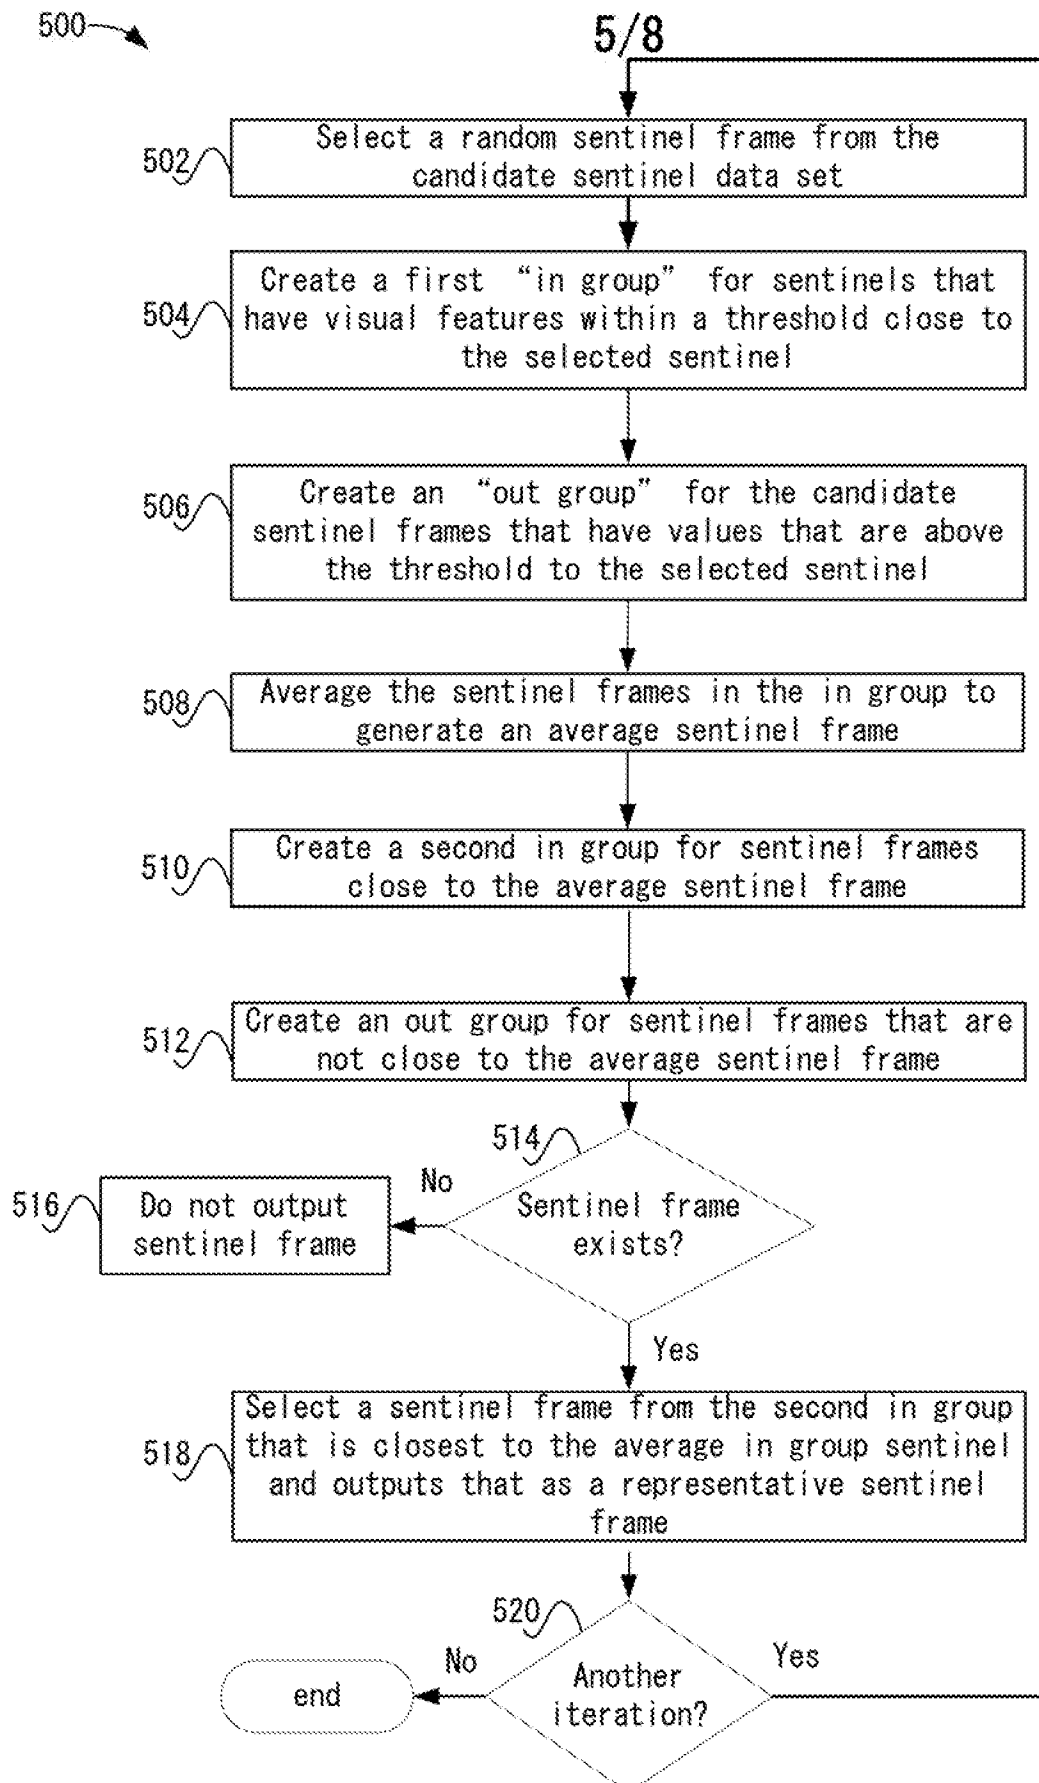


FIG. 5

6/8




FIG. 6

7/8

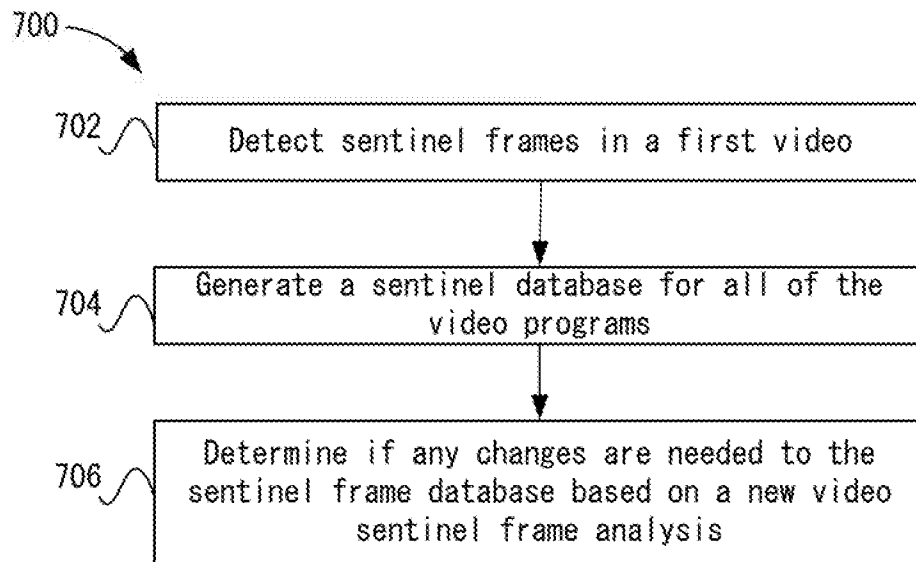


FIG. 7




FIG. 8

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/032117

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06K9/00
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06K

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2016/037232 A1 (HU SI YING DIANA [US] ET AL) 4 February 2016 (2016-02-04)	1-11,20
Y	abstract paragraph [0041] - paragraph [0047] paragraph [0052] - paragraph [0054] -----	12-19
Y	US 2004/170321 A1 (GONG YIHONG [US] ET AL) 2 September 2004 (2004-09-02) abstract paragraph [0067] - paragraph [0079] -----	12-19



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

25 July 2017

Date of mailing of the international search report

01/08/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Angelopoulou, Maria

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2017/032117

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016037232	A1	04-02-2016	NONE

US 2004170321	A1	02-09-2004	JP 3568117 B2 22-09-2004
			JP 2001155169 A 08-06-2001
			US 7016540 B1 21-03-2006
			US 2004170321 A1 02-09-2004
			US 2006280365 A1 14-12-2006
