



- (51) International Patent Classification:
G06K 9/00 (2006.01) *G06F 17/30* (2006.01)
- (21) International Application Number:
PCT/US2015/033662
- (22) International Filing Date:
2 June 2015 (02.06.2015)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
14/302,155 11 June 2014 (11.06.2014) US
- (71) Applicant: **ARRIS ENTERPRISES, INC.** [US/US];
3871 Lakefield Drive, Suwanee, Georgia 30024 (US).
- (72) Inventors: **FONSECA, Benedito J., Jr.**; 601 Lowden Avenue, Glen Ellyn, Illinois 60137 (US). **ISHTIAQ, Faisal**; 5415 Sheridan, Chicago, Illinois 60640 (US). **LI, Renxiang**; 290 Stone Avenue, Lake Zurich, Illinois 60047 (US). **EMEOTT, Stephen P.**; 5608 Silentbrook Lane, Rolling Meadows, Illinois 60008 (US). **SMITH, Alfonso Martinez**; 650 Oakview Drive, Algonquin, Illinois 60102 (US). **BRASKICH, Anthony J.**; 227 E. Palatine Road, Palatine, Illinois 60067 (US).
- (74) Agents: **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, 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, KN, KP, KR, 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.
- (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))

(54) Title: SELECTION OF THUMBNAILS FOR VIDEO SEGMENTS




FIG. 1

(57) Abstract: A method of identifying a representative image of a video stream is provided. Similarity between video frames of a primary video stream relative to video frames of a different secondary video stream having similar content is evaluated and a video frame from the primary video stream having a greatest extent of similarity relative to a video frame of the secondary video stream is identified. The identified video frame is selected as an image representative of the primary video stream and may be used as an informative thumbnail image for the primary video stream. A video processing electronic device and at least one non-transitory computer readable storage medium having computer program instructions stored thereon for performing the method are also provided.

SELECTION OF THUMBNAILS FOR VIDEO SEGMENTS

BACKGROUND

[0001] Broadband network operators, such as multiple system operators (MSOs), distribute and deliver services such as video, audio, and multimedia content to subscribers or end-users. For example, a broadband cable network MSO may utilize resources for transmitting digital video as linear (i.e., scheduled) services or as non-linear services enabling viewers to retrieve audiovisual contents at any time independent from linear (i.e., scheduled) broadcast services.

[0002] Some non-linear content may be of a type readily capable of being presented to users in distinct segments thereby permitting the user to view any of the segments without respect to other segments and in any desired sequence. By way of example, a newscast represents a type of video that typically includes numerous distinct segments which each may be directed to a separate news worthy event or a separate topic, for instance, related to weather, sports, entertainment, or like subject matter. Thus, if the newscast is segmented into distinct parts, the viewer may consume the newscast as a non-linear service after the original broadcast of the newscast and view only the segments of the newscast of particular interest to the viewer and in any desired sequence.

[0003] When a video asset is segmented and provided as non-linear content, for instance, by an MSO as discussed above, it is typically desirable to present the viewer with separate links to the segments via a graphical user interface that includes images, typically referred to as a “thumbnails”, representative of the subject matter to which the segments are directed. Accordingly, upon segmenting a video asset, it is desirable to

assign thumbnails or representative images informative of the subject matter content of each segment.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] Various features of the embodiments described in the following detailed description can be more fully appreciated when considered with reference to the accompanying figures, wherein the same numbers refer to the same elements.

[0005] FIG. 1 is an exemplary image of a graphical user interface (GUI) having thumbnails in accordance with an embodiment.

[0006] FIG. 2 is a diagram representing a process for segmenting a first video stream in accordance with an embodiment.

[0007] FIG. 3 is a diagram representing a process for extracting features from video frames of the first video stream in accordance with an embodiment.

[0008] FIG. 4 is a diagram representing a process for obtaining a second video stream having overlapping content with the first video stream in accordance with an embodiment.

[0009] FIG. 5 is a diagram representing a process for extracting features from video frames of the second video stream in accordance with an embodiment.

[0010] FIG. 6 is a diagram representing a process for determining a similarity metric between a first selected video frame of the first video stream and video frames of the second video stream in accordance with an embodiment.

[0011] FIG. 7 is a diagram representing a process for determining a similarity metric between a second selected video frame of the first video stream and video frames of the second video stream in accordance with an embodiment.

[0012] FIG. 8 is a diagram representing a process for determining a similarity metric between a third selected video frame of the first video stream and video frames of the second video stream in accordance with an embodiment.

[0013] FIG. 9 is a diagram representing selection of a representative image for use as a thumbnail for a first segment of the first video stream in accordance with an embodiment.

[0014] FIG. 10 is a diagram of system architecture including a thumbnail selector in accordance with an embodiment.

DETAILED DESCRIPTION

[0015] For simplicity and illustrative purposes, the principles of the embodiments are described by referring mainly to examples thereof. In the following description, numerous specific details are set forth in order to provide a thorough understanding of the embodiments. It will be apparent however, to one of ordinary skill in the art, that the embodiments may be practiced without limitation to these specific details. In some instances, well known methods and structures have not been described in detail so as not to unnecessarily obscure the embodiments.

[0016] According to an embodiment, a graphical user interface (GUI) or like interface is used to present links to non-linear content provided in the form of separate

and distinct video segments. With such a service, the user is permitted to select and view any of the video segments without respect to the other available video segments and in any sequence. As suggested above, a newscast provides an example of a video stream that is of a type that can be readily divided into separate video segments with each segment being directed to a different news event, weather, sports, entertainment, or like subject matter. Of course, a newscast is only being disclosed for purposes of example, and embodiments disclosed herein are equally applicable to any video stream that is of a type that may be provided to users in a segmented format. By providing such a video stream in segmented form, the viewer is able to consume the content as non-linear content with control over which video segment or segments of particular interest to the viewer are played in any desired sequence.

[0017] FIG. 1 provides an example of a GUI 10, for instance, for a newscast. The main display area 12 of the GUI 10 may provide an indication of the video stream, such as the name of the particular newscast or its source. Different segments may be represented by thumbnails 14, 16, 18, 20 and 22 in the GUI 10. Thus, by selecting one of the thumbnails, the user can cause the video, audio and/or other multimedia content of the corresponding segment to be displayed or begin playing. It is preferable that the image displayed in each thumbnail 14, 16, 18, 20 and 22 is in some way representative of the subject matter content of its corresponding video segment to enable the user to readily ascertain the particular subject matter content of each segment based solely on viewing the thumbnail images and without having to actually begin to view the video segment.

[0018] If thumbnails are not manually pre-selected for each segment, which entails a highly labor intensive process considering the amount of content and variation of subject matter throughout video assets, it may be necessary to automatically generate thumbnails. As an example of one possible approach, the nth video frame of each video segment may be automatically selected as an image for a thumbnail. However, with this approach, the selected thumbnail may not necessarily represent the actual content and may not be informative as to the subject matter of the segment. For instance, each video segment may include multiple scenes and one or more scene cuts, and it may be difficult to determine which scene within a video segment is informative of, and best represents the content of, the video segment. By way of example, an image or thumbnail of a news anchor seated behind a news desk may fail to provide any relevant information as to the subject of a news segment.

[0019] According to an embodiment, a method of automatically identifying a representative image of a video segment of a primary video stream or asset utilizes at least one secondary video stream or asset for purposes of making a comparison. The secondary video stream is a video stream that is necessarily different from the first video stream but which may contain related, overlapping, or similar content. For example, if the primary video stream or asset is a newscast, the secondary video stream may be a similar newscast obtained from another source or channel or may be a newscast from the same producer, source or channel that may have preceded the primary video stream by a few hours or other acceptable time period. Thus, similar news topics may be presented in

the primary and secondary video streams or assets; however, not necessarily in an identical format or sequence.

[0020] In the above referenced embodiment, the process may include the computation or development of some form of similarity metric for purposes of making a comparison between the primary and secondary video assets. For example, features, such as image descriptors, generated at a selected time point or frame time of the primary digital video stream may be compared to similar features generated at numerous time points or frame times of the secondary digital video stream for purposes of determining the similarity of the features. Such a process may be repeated for various time points or frame times of the primary video stream for purposes of computing the similarity metric corresponding to each time point, frame time, or video frame. The computed similarity metrics can then be used to automatically select a video frame or frames of the primary video stream having the highest similarity to a video frame or frames of the secondary video stream for use as a thumbnail to present to a viewer in a GUI. By way of example, video frames within two different newscasts that have high similarity according to the above process should represent video frames that will be most relevant to a viewer in promptly ascertaining an understanding of the nature of the content of the corresponding video segment.

[0021] With respect to the above discussed newscast example, the process of the embodiment may include a MSO receiving a newscast of a first video asset and segmenting the video asset into separate and distinct segments. For each segment, the MSO may extract audio, video, and /or text features or descriptors at various time points

of the segment. In addition, the MSO accesses a second newscast from another source such as a newscast from a different channel or from the same producer at an earlier time. Such second newscast may have been aired in the same day or within a few hours from the first newscast so that the newscasts are likely to contain overlapping subject matter. The MSO extracts text, audio and/or video features or descriptors at various time points from the second newscast. The MSO then evaluates similarity of the extracted features by computing a similarity metric. For instance, the text, audio and/or video features at each of the time points of a segment of the first newscast may be compared to the text, audio and/or video features at time points of the second newscast. The time point in the first video stream that provides a highest similarity metric with any time point in the second video stream is automatically selected as the thumbnail representative image for the segment of the first newscast. These steps can be repeated for each of the segments of the first newscast to generate an array of thumbnails for the segmented first newscast.

[0022] FIGs. 2-9 provide an example of the above referenced method applied to a newscast asset. In FIG. 2, a first newscast 30 is provided from a source 32 to a service provider, or MSO, 34 in the form of a digital video stream that may include audio, multimedia, closed captions, metadata, and the like. The service provider 34 segments the various sections of the first newscast 30 such that the segments, such as video segments 36, 38 and 40, may be consumed by end users in a non-linear fashion via use of a GUI, such as shown in FIG. 1. As shown in the example provided by FIG. 2, “segment 1” (reference numeral 36) may be directed to a new event including a president’s speech, “segment 2” (reference numeral 38) may be an advertisement, and “segment 3”

(reference numeral 40) may be a news event with respect to a fire at a warehouse. The service provider 34 may use text, audio and/or video features or the like to segment the first newscast 30 and determine where each segment begins and ends. In some embodiments, the segments' contents are not necessarily labeled; in other embodiments, the segments' contents may have been labeled by a system that provides automatic topic summarization.

[0023] The video timeline 42 in FIG. 3 shows a plurality of time points 44 along the digital video stream of the segment 36 of the first newscast 30. The time points 44 may represent frame times or the first frame time of a group of consecutive frames. For each time point, or for each selected time point, on the timeline 42, a set of features can be extracted. For instance, 1 to N_v video features (v_f) may be extracted and 1 to N_a audio features (a_f) may be extracted for each time point subject to evaluation. Text features may also be extracted, such as text appearing directly within the video frame or text contained in closed caption data or metadata corresponding to the time point or corresponding video frame.

[0024] As an example, the features may be visual descriptors or image descriptors of the visual features of the contents in images, video frames, or videos. The features may be used to describe elementary characteristics such as shape, color, texture or motion. Color Layout Descriptors (CLDs) and Edge Histogram Descriptors (EHDs) provide two examples of features that can be used. A set of Color Layout Descriptors (CLDs) can be used to capture spatial distribution of color in an image or frame of video, and a set of Edge Histogram Descriptors (EHDs) can be used to capture

the texture of an image. Accordingly, when such features or descriptors of two similar images are compared, the comparison will produce a similarity metric indicating high similarity. In contrast, when such features or descriptors of two significantly different images are compared, the comparison will produce a similarity metric indicated low or no similarity.

[0025] There are many ways to compute the similarity metric. In one embodiment, the similarity metric is obtained by computing the inner product between two vectors containing video features. For instance, if the video feature corresponds to color layout descriptors, the inner product between the vector of coefficients in the color layout descriptor of a first image and the corresponding vector of coefficients of a second image would provide a similarity metric. In another embodiment, the similarity metric is obtained from the weighted sum of the square of the difference between the coefficients in the color layout descriptor of a first image and the corresponding coefficients of the color layout descriptor of a second image.

[0026] With respect to the newscast example, the use of unique graphics (banners, framing, etc.) added by producers, for instance, of different channels, may need to be considered when making such comparisons as the graphics themselves may alter the outcome. For at least this purpose, the determination and use of audio features or descriptors or text features may be particularly useful when the content, such as a news piece, is about or contains a particular object, person, or event, such as a president's speech, an official's statement, explosions, crowd noise, sirens, or the like.

[0027] Each video frame of the video stream of the segment 36 may be evaluated and considered as a candidate for selection as a thumbnail for the segment 36. According to one embodiment, each video frame is subject to evaluation as a candidate, and according to other embodiments, only selected video frames are subject to evaluation as a candidate. For example, the number of video frames subject to evaluation as a candidate may be reduced by restricting candidates to video frames occurring only at a scene cut within the segment. Other criteria may also be used, such as video frames occurring at preset intervals along timeline 42 or the like. Thus, all video frames may be considered and evaluated as candidates for use as a thumbnail or only video frames located at predetermined intervals, at scene cuts, or the like may be considered and evaluated as candidates.

[0028] In FIG. 4, the service provider 34 accesses video streams and assets from other sources 46 and 48 that are not identical to the first newscast 30 but which may contain similar or overlapping content. For example, a second newscast 50 and a third newscast 52 that was produced relatively close in time to the first newscast 30 may be obtained. The second or third newscasts 50 and 52 may be from a preceding day or from earlier in the same day, for instance, within three hours or like predetermined time period. As shown in FIG. 4, preferably the second and/or third newscasts 50 and 52 should have similar or overlapping content, although not necessarily provided in the same sequence as the first newscast 30.

[0029] As shown in FIG. 5, a video timeline 54 shows a plurality of time points 56 along the video stream of the second newscast 50. The time points 56 may represent

frame times or the first frame time of a group of consecutive frames. For each time point or for selected time points along the timeline 54, a set of features or descriptors can be extracted as accomplished for the first newscast 30. For instance, 1 to N_v video features (v_f) can be extracted and 1 to N_a audio features (a_f) may be extracted for each time point being evaluated. Text features may also be extracted, such as text appearing within the video frame or text within closed caption data or metadata corresponding to the video frame. Each video frame of the entire video stream of the second newscast 50 may be evaluated or evaluation may be limited to video frames determined to be scene cuts or the like. Thus, all video frames may be used or only a subset of selected video frames may be considered in computing similarity metrics.

[0030] FIG. 6 shows the timeline 42 of the video segment 36 of the first newscast 30 relative to the timeline 54 of the entire second newscast 50 or relevant sections thereof. A similarity metric is computed for each time point, frame time, or video frame of the video segment 36 considered as a candidate for use as a thumbnail relative to each time point, frame time, or video frame of the second newscast 50. The purpose of the similarity metric is to identify a video frame in the segment 36 that closely matches a video frame in the second newscast 50 so that such a video frame of the segment 36 can be used as a representative image or thumbnail for the segment 36 of the first newscast 30.

[0031] By way of example, the set of features determined and corresponding to video frame 58 of the video segment 36 of the first newscast 30 are separately compared to each of the sets of features determined for each or selected ones of the time points or

video frames of the second newscast 50. FIG. 7 shows this process repeated for another time point of the segment 36 corresponding to video frame 60, and FIG. 8 shows this process repeated for a further time point of the segment 36 corresponding to video frame 62. FIG. 9 shows the result of all of the above referenced comparisons. The time point corresponding to video frame 62 is determined to have the greatest amount of similarity, based on a comparison of extracted features, with a time point corresponding to video frame 64 of the second newscast 50. Here, segment 36 of the first newscast 30 is directed to a president's speech. The second newscast 50 also covered the president's speech in a second segment thereof. Thus, the video frames 62 and 64 showing the president giving a speech provides the highest similarity metric computed between frames of segment 36 relative to frames of the second newscast 50 and is therefore automatically selected for use as a thumbnail representative of the content of video segment 36 of the first newscast 30.

[0032] The process shown in FIGs. 6-9 may be repeated for each of the segments of the first newscast 30, such as segments 38 and 40. Thus, a different thumbnail may be automatically selected for each of the segments 36, 38 and 40 of the first newscast 30, and each of the selected thumbnails should be informative of the content each segment. The selected representative images may be used as thumbnails for display in a user interface such as shown in FIG. 1 and displayed to a subscriber of the service provider 34 for use in selecting and viewing any of the segments in the form of non-linear content.

[0033] An embodiment of a system 66 of electronic video processing apparatus for automatically performing the above method is shown in FIG. 10. The system 66 may

include one or more segment generators 68 for receiving first, second, third, etc. video streams. Each segment generator 68 segments the video streams including applicable metadata such as closed caption information or the like. The various segments of the various streams may be stored in a storage unit 70. A thumbnail selector 72 including one or more processors may be configured to extract features from the segments of a primary video stream and make comparisons with secondary video streams to determine video frames of highest similarity so that thumbnails can automatically be selected for each segment of a primary video stream or asset as discussed above. Thus, features and/or descriptors and similarity metrics can be computed between each time point or interval of the first video sequence and all of the time points or intervals of the second video sequence. The automatically selected thumbnails can then be used in a user interface 74 to provide informative representative images to potential viewers of any of the video segments.

[0034] Accordingly, a thumbnail can be generated automatically for a first video stream, or for segments thereof, based on a comparison of features of video frames with a secondary video stream of similar nature and content. In this manner, the most relevant video frame can be identified based on the comparison of video frames and features computed therefrom and an informative thumbnail can be provided.

[0035] Various modifications can be made to the above referenced embodiments. For example, the time points in the primary video representing candidate video frames subjected to evaluation and consideration as thumbnail images may be selected based on frame times, regularly spaced time points in the video segment (e.g., every 5 seconds of

video), time points in the video in which a scene cut appears and can be detected, time points in the video in which audio characteristics change, time points in the video in which the short-term statistics of any combination of text, audio, and video features change by a predetermined threshold amount, or on any basis.

[0036] For example, if limiting the number of candidates to be subject to evaluation is desired, for instance, for purposes of reducing processing time or requirements of the video stream, the time points of the video streams of the first segment under analysis and the various time points of the second or other video asset may be trimmed. For instance, as suggested above, the number of candidate time points for evaluation as thumbnails may be restricted solely to time points in the video streams which involve scene cuts.

[0037] Another alternative is to calculate a similarity metric of text, audio and video features between a desired segment of the first video stream and text, audio and video features from other segments of the first video stream (i.e., of the same asset). This process step can be used for purposes of excluding many time points in the video stream under analysis in which, for instance, a newscaster appears since the same newscaster may appear in many points of the newscast and may not be informative of the subject matter of the segment. With this step, similar video frames occurring within the same video asset may be eliminated from being evaluated and considered as a candidate for use as a thumbnail.

[0038] According to another embodiment, a set of filters is applied to the video frames of the segment of the first video stream under analysis for purposes of excluding

some of the frames from further evaluation. For example, face recognition techniques and software could be used to exclude frames containing faces from known news anchors and reporters. Alternatively, speech recognition techniques and software may be used to exclude portions of the video segments that contain speech from known news anchors and reporters. Still further, object detection techniques and software may be used to enhance the selection of thumbnails. For instance, if text or keywords in the closed captions of the segment being evaluated concerns airplanes or some other object, only the frames of the video segment of the first video stream that contain one or more airplanes would be considered as candidates for evaluation for thumbnail selection. Furthermore, only frames of the second video segment that contain one or more airplanes or other object may be used to compute the similarity metric.

[0039] Yet a further embodiment may utilize face detection techniques and software to enhance the selection of thumbnails. For example, if text analysis of the closed captions of the video segment identifies that the news segment is about a person, athlete, or celebrity, then only video frames of the first video segment that contain one or more of the identified faces would be considered as candidates for evaluation as thumbnails. Furthermore, only frames of the second video segment that contain one or more of the identified faces would be used to compute the similarity metric.

[0040] Similarity metrics of text, audio, and video features may be calculated between a desired segment of a primary video stream and text, audio, and/or video features from other segments of previous assets of the same channel or producer. Here, if an image appears in several days of a newscast, it is likely that this image does not

represent images from news of the current day. Since some news pieces may be relevant for a duration of several days, even weeks or months, when using this particular technique, a candidate time point would only be eliminated from further evaluation if it has a relatively high similarity with previous images of most or all of the previous newscasts for a period of time. Various modifications can be made to these embodiments with respect to reducing the number of video frames that are ultimately subject to evaluation as candidates for use as a thumbnail.

[0041] In accordance to another embodiment, each segment of the primary and secondary video streams may be subject to initial partitioning into parts. One part may be video including an anchor or reporter and another part may be video in which an anchor or reporter is not detected. Thereafter, evaluation of candidate video frames may proceed only for video frames in the part of the video in which an anchor or reporter is not detected in the primary video stream and the video frames of the secondary video stream.

[0042] A further embodiment with respect to evaluating and comparing video frames for similarity may require the similarity metric to be computed based on a collective similarity metric involving use of a group of two or more consecutive frames of a segment of a first video stream relative to a group of two or more consecutive frames of a second video stream. Thus, evaluation may involve a similarity metric determined based on a comparison of features of an individual video frame relative to other individual video frames, or evaluation may involve a collective similarity metric obtained for a number of consecutive video frames. In the latter case, a thumbnail for a video segment of the first video stream may be selected as the nth frame of the group of

consecutive video frames determined to have the greatest similarity metric. As another alternative, the thumbnail could be provided as a short video clip of the selected group of consecutive video frames instead of a static image.

[0043] As a further option, a single set of features may be extracted from the video of the group of two or more consecutive frames of the segment of the first video stream and the same set of features may be extracted from the video of the group of two or more consecutive frames of the second video stream. The similarity metric can then be computed based on a comparison between the two sets of extracted features.

[0044] According to another embodiment, if a second newscast is used to generate a thumbnail for a segment of a first newscast, the first newscast may be used to generate a thumbnail for a corresponding segment of the second newscast. Thus, the first and second video stream would simultaneously be subject to evaluation and thumbnail selection process. In addition, as segments from different video streams are determined to refer to the same content, links between such segments can be automatically generated and provided to the viewer in the user interface.

[0045] According to another embodiment, keywords obtained from a closed-caption stream or other metadata provided with the video streams can be utilized in the thumbnail selection process. For example, keywords in a segment of the primary video stream could be identified and used in searches for secondary videos containing like keywords. Thus, video streams on the Internet, YouTube, or like sources could be searched and identified as an asset directed to similar content. Thus, a similarity metric

as discussed above could be computed between video frames of a video segment of the first video stream and the YouTube or like content obtained in the keyword search.

[0046] Thus, any of the embodiments disclosed above may also include process steps for selecting other secondary video stream sources for comparison to the primary or first video stream. For example, these steps may include verifying the genre of the first asset (i.e., the asset for which a thumbnail is needed), verifying the genre of any other available asset, and selecting one or more of the other available assets based on the asset being of the same genre as the first asset. According to some embodiments, one or more assets with a different genre may be selected if a sufficiently high similarity metric could not be contained with the initially selected other assets.

[0047] As a further alternative, keywords from the closed captions of the segment in the first asset (i.e., the asset for which a thumbnail is needed) may be identified and extracted, the closed captions of other available assets can be identified and extracted, and selection of assets having keywords in common with the first asset can be selected. In some embodiments, only closed captions/keywords of assets that have the same genre as the first asset may be searched.

[0048] While most of the above examples primarily focus on automatically identifying thumbnails for newscasts, this is only provided by way of example. The above methods and systems can be applied to other video streams. For example, the above techniques can be applied to video stream that have different segments including talk shows, sports shows, late-night shows, variety shows, music shows, and the like.

[0049] The above referenced system, apparatus, and processing units may include various processors, microprocessors, controllers, chips, disk drives, and like electronic components, modules, equipment, resources, servers, and the like for carrying out the above methods and may physically be provided on a circuit board or within another electronic device. It will be apparent to one of ordinary skill in the art that the processors, controllers, modules, and other components may be implemented as electronic components, software, hardware or a combination of hardware and software.

[0050] For example, at least one non-transitory computer readable storage medium having computer program instructions stored thereon that, when executed by at least one processor, can cause the at least one processor to automatically select thumbnails for a video stream as described in the above embodiments.

[0051] In the foregoing specification, specific embodiments have been described. However, one of ordinary skill in the art appreciates that various modifications and changes can be made without departing from the scope of the embodiments as set forth in the claims below. Accordingly, the specification and figures are to be regarded in an illustrative rather than a restrictive sense, and all such modifications are intended to be included within the scope of the embodiments.

We claim:

1. A method of identifying a representative image of a video stream, comprising the steps of:

evaluating similarity between video frames of a primary video stream

relative to video frames of a different secondary video stream

having similar content;

identifying a video frame from the primary video stream having greatest

similarity relative to a video frame of the secondary video stream;

and

selecting the video frame identified in said identifying step as an image

representative of the primary video stream for use as thumbnail for

the primary video stream.

2. A method according to claim 1, wherein said evaluating step includes a step of computing a similarity metric between video frames of the primary video stream and video frames of the secondary video stream.

3. A method according to claim 2, wherein said identifying step includes selecting a video frame of the primary video stream corresponding to a similarity metric computed during said computing step indicating greatest similarity.

4. A method according to claim 2, wherein said computing step includes extracting features from said video frames of said primary video stream and video frames of said secondary video stream and comparing said features.

5. A method according to claim 4, wherein said features include at least one of video, text, and audio features.

6. A method according to claim 2, wherein the video frames of the primary video stream used during said computing step are selected from a group consisting of all video frames, video frames at predetermined intervals, video frames corresponding to scene cuts, video frames in which at least one of video, text and audio features change.

7. A method according to claim 6, wherein the video frames of the secondary video stream used during said computing step are selected from a group consisting of all video frames, and video frames corresponding to scene cuts.

8. A method according to claim 2, wherein the similarity metric is computed between a group of two or more consecutive video frames of the primary video stream and two or more consecutive frames of the secondary video stream.

9. A method according to claim 1, further comprising the step of causing the identified video frame to be displayed as a thumbnail for representing the primary video stream in a user interface.

10. A method according to claim 1, further comprising the step of dividing the primary video stream into a plurality of video segments and repeating said evaluating, identifying and selecting steps for each of the plurality of video segments such that a different representative image is selected for each of the plurality of video segments.

11. A method according to claim 10, wherein the secondary video stream includes a plurality of segments, and wherein the representative image selected during said selecting step for one of the plurality of video segments of the primary video segment is selected as a representative image of a corresponding segment in the secondary video segment.

12. A method according to claim 10, further comprising the steps of:
calculating a similarity metric between video frames of one of the plurality
of video segments of the primary video stream and video frames of
at least a different one of the plurality of video segments of the
primary video stream before said evaluating step; and

excluding video frames of the primary video stream from said evaluating step that have similarity metrics, determined during said calculating step, to exceed a preset threshold.

13. A method according to claim 1, wherein the primary video stream corresponds to a newscast and the secondary video stream corresponds to an earlier newscast relative to the newscast of the primary video stream.

14. A method according to claim 13, further comprising the steps of:
identifying video frames including news anchors or reporters in the
primary and secondary video streams before said evaluating step;
and
excluding the video frames including news anchors or reporters from said
evaluating step.

15. A method according to claim 1, further comprising the step of obtaining the secondary video stream by searching keywords from a closed caption stream of the primary video stream.

16. A method according to claim 1, further comprising the steps of:
obtaining keywords from a closed caption stream of the primary video stream;

identifying video frames in the primary and secondary video streams
including objects or persons corresponding to the keywords; and
performing said evaluating step only with video frames of the primary and
secondary video streams including objects or persons
corresponding to the keywords.

17. A video processing electronic device for automatically identifying a
representative image of a video stream, comprising at least one processing unit
configured to:

evaluate similarity between video frames of a primary video stream
relative to video frames of a different secondary video stream
having similar content;
identify a video frame from the primary video stream having greatest
similarity relative to a video frame of the secondary video stream;
and
selecting the identified video frame as an image representative of the
primary video stream.

18. A video processing electronic device according to claim 17, wherein said at
least one processing unit is configured to compute a similarity metric between video
frames of the primary video stream and video frames of the secondary video stream and


to select a video frame of the primary video stream corresponding to a highest similarity metric computed.

19. A video processing electronic device according to claim 18, wherein said at least one processing unit is configured to compute the similarity metric by extracting features from said video frames of said primary video stream and video frames of said secondary video stream and to compare said features, wherein said features include at least one of video, text, and audio features.

20. At least one non-transitory computer readable storage medium having computer program instructions stored thereon that, when executed by at least one processor, cause the at least one processor to automatically identify a representative image of a video stream by performing the following operations:

- evaluate similarity between video frames of a primary video stream
 - relative to video frames of a different secondary video stream
 - having similar content;
- identify a video frame from the primary video stream having greatest
 - similarity relative to a video frame of the secondary video stream;
 - and
- selecting the identified video frame as an image representative of the
 - primary video stream.

1/10






FIG. 2

3/10




FIG. 3




FIG. 4






FIG. 5


6/10



7/10



8/10



9/10




FIG. 9

10/10




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/033662

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06K9/00 G06F17/30
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 G06F

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, COMPENDEX, INSPEC, 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	Howard Wactlar: "Lessons Learned from Building a Terabyte Digital Video Library", Computer, 1 February 1999 (1999-02-01), XP055210671, Retrieved from the Internet: URL:http://www.cs.cmu.edu/~wactlar/IEEEComputer Feb99.pdf [retrieved on 2015-09-01] page 67-page 68, Figure 1 ----- -/--	1-20



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

2 September 2015

Date of mailing of the international search report

14/09/2015

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

Versluis, Anton

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/033662

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	FLICKNER M ET AL: "QUERY BY IMAGE AND VIDEO CONTENT: THE QBIC SYSTEM", COMPUTER, IEEE, US, vol. 28, no. 9, 1 September 1995 (1995-09-01), pages 23-32, XP000673841, ISSN: 0018-9162, DOI: 10.1109/2.410146 Figure 7 and its caption -----	1
A	US 2014/099034 A1 (RAFATI SHAHRZAD [CA] ET AL) 10 April 2014 (2014-04-10) -----	1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2015/033662

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014099034 A1	10-04-2014	US 2014099034 A1	10-04-2014
		WO 2014056112 A1	17-04-2014
