

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2016291690 B2**

(54) Title
Prediction of future views of video segments to optimize system resource utilization

(51) International Patent Classification(s)
H04N 21/462 (2011.01)

(21) Application No: **2016291690** (22) Date of Filing: **2016.07.15**

(87) WIPO No: **WO17/011792**

(30) Priority Data

(31) Number	(32) Date	(33) Country
62/193,331	2015.07.16	US

(43) Publication Date: **2017.01.19**

(44) Accepted Journal Date: **2020.08.27**

(71) Applicant(s)
Inscape Data, Inc.

(72) Inventor(s)
Neumeier, Zeev; Collette, Michael

(74) Agent / Attorney
Davies Collison Cave Pty Ltd, Level 14 255 Elizabeth St, Sydney, NSW, 2000, AU

(56) Related Art
US 2014/0201772 A1
LE SCOUARNEC, N. ET AL.: "Cache policies for cloud-based systems: To keep or not to keep", 2014 IEEE 7th International Conference on Cloud Computing, 27 June 2014, pages 1-8
US 2014/0025837 A1



- (51) International Patent Classification:
G06F 17/30 (2006.01) H04N 21/462 (2011.01)
- (21) International Application Number:
PCT/US2016/042611
- (22) International Filing Date:
15 July 2016 (15.07.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/193,331 16 July 2015 (16.07.2015) US
- (71) Applicant: VIZIO INSCAPE TECHNOLOGIES, LLC
[US/US]; 39 Tesla, Irvine, California 92618 (US).
- (72) Inventors: NEUMEIER, Zeev; 39 Tesla, Irvine, California 92618 (US). COLLETTE, Michael; 19 Tesla, Irvine, California 92618 (US).
- (74) Agents: MORSE, Kyle B. et al.; Kilpatrick Townsend & Stockton LLP, 1100 Peach Street, Suite 2800, Atlanta, Georgia 30309 (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: PREDICTION OF FUTURE VIEWS OF VIDEO SEGMENTS TO OPTIMIZE SYSTEM RESOURCE UTILIZATION

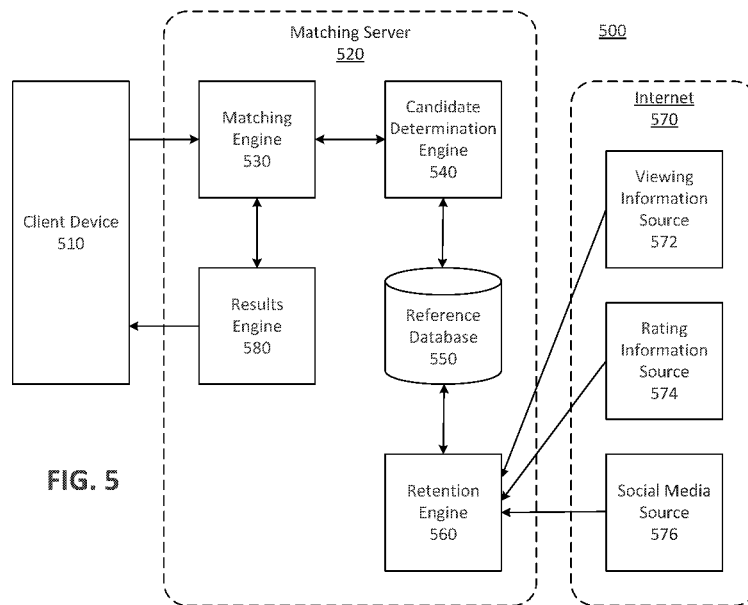


FIG. 5

(57) Abstract: Provided are devices, computer-program products, and methods for improved management of system resources in a matching system. For example, examples can increase the efficiency of system resource utilization by managing the duration that data related to video segments are retained based on data that takes into account an identified popularity of a video segment. The identified popularity can be determined by algorithms that take into account numbers of viewers who watched the video segment, ratings of the video segment, metrics derived from remote sources, or any other factor that can indicate likelihood that the video segment will be viewed.

WO 2017/011792 A1

**PREDICTION OF FUTURE VIEWS OF VIDEO SEGMENTS TO
OPTIMIZE SYSTEM RESOURCE UTILIZATION**

5

CROSS-REFERENCES TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 62/193,331, filed on July 16, 2015, which is hereby incorporated by reference in its entirety. This application also incorporates by reference the following in their entirety: U.S. Patent Application Serial No. 10 US 14/089,003, which was filed on November 25, 2013, U.S. Provisional Application No. 61/182,334, which was filed on May 29, 2009, U.S. Provisional Application No. 61/290,714, which was filed on December 29, 2009, U.S. Application No. 12/788,748, which was filed on May 27, 2015, and U.S. Application No. 12/788,721, which was filed on May 27, 2015.

15

FIELD

[0002] The present disclosure relates to improving management of system resources used for recognition of content displayed by a media system (e.g., a television system, a computer system, or other electronic device capable of connecting to the Internet). In some examples, various techniques and systems are provided for determining a likelihood that content will be displayed 20 on a media system in a future period of time.

BACKGROUND

[0003] A media system (e.g., a television system, a computer system, or other electronic device capable of connecting to the Internet) can display content for a user viewing the media 25 system. In some examples, the content displayed on the media system might be unknown to the media system, and any application run on the media system. For examples, the media system might display whatever the media system is given. In such examples, an identification of the content might not be sent to the media system. In other examples, the identification of the content might have been lost between when the remote system sent the content and when the

media system received the content. By not knowing what is being displayed on the media system, opportunities to show similar or relevant additional content with the content on the media system becomes difficult.

[0004] One solution to identify content displayed on a media system is to use a matching system, which matches content displayed on the media system with reference content. However, in order to be able to match content, the matching system must include a large number of reference content. In addition, not only must the matching system have a large number of reference content, the matching system also must have the right reference content. Therefore, there is a need in the art for managing reference content for a matching system.

10

SUMMARY

[0005] Provided are devices, computer-program products, and methods for improved management of system resources in a matching system. For example, examples can increase the efficiency of system resource utilization by managing the duration that data related to video segments are retained based on data that takes into account an identified popularity of a video segment. The identified popularity can be determined by algorithms that take into account numbers of viewers who watched the video segment, ratings of the video segment, metrics derived from social media, or any other factor that can indicate likelihood that the video segment will be viewed.

15

[0006] In some implementations, a device, computer-program product, and method for improving management of system resources in a matching system is provided. For example, a method can include receiving a reference data set associated with a media segment and storing the reference data set in a reference database. In some examples, the reference data set can be received by a server. The server can be configured to identify an unidentified media segment by matching an unidentified data set with the reference data set. In such examples, the unidentified data set can be associated with the unidentified media segment. In some examples, a reference data set can include pixel data or audio data for a frame of the media segment.

20

25

[0007] The method can further include receiving popularity data indicating a popularity of the media segment. In some examples, the popularity data includes at least one or more of viewing

information of the media segment, rating information of the media segment, or information associated with a posted message on a remote source. In such examples, the posted message can be associated with the media segment. In some examples, the information associated with the posted message on the remote source can include at least one or more of a number of posted messages on the remote source, a number of the posted messages that are positive, a number of the posted messages that are negative, a number of positive indications to the posted messages, a number of reposts of the posted message, or a number of views of the posted messages. In such examples, the posted messages are associated with the media segment. In some examples, the viewing information can include a number of times the media segment has been identified by the server.

[0008] The method can further include determining a deletion time for the reference data set. In some examples, the deletion time can be determined using the popularity data. The method can further include deleting the reference data set from the reference database. In some examples, the reference data set can be deleted after the deletion time has expired. In some examples, determining a deletion time can include computing an average of two or more values of the popularity data. In such examples, each value of the popularity data can be associated with a weight.

[0009] In some implementations, the method can further include adding the reference data set to a backup database when the reference data set is deleted from the reference database. In such implementations, the server can be configured to search the backup database after the reference database for a data set matching the unidentified data set.

[0009A] In another example aspect, there is provided a computer-implemented method comprising: receiving, by a server, a reference data set associated with a media segment, wherein the server is configured to identify an unidentified media segment by matching an unidentified data set with the reference data set, and wherein the unidentified data set is associated with the unidentified media segment; storing the reference data set in a reference database; receiving popularity data indicating a popularity of the media segment, wherein the popularity data includes values associated with at least one or more of viewing information of the media segment, rating information of the media segment, or information associated with a posted message on a remote source, and wherein the posted message is associated with the media segment; determining a deletion time for the reference data set, wherein the deletion time is determined by calculating a representative value of two or more values of the popularity data,

the representative value being one of an average, a mean, or a median of the two or more of the values of the popularity data; and deleting the reference data set from the reference database, wherein the reference data set is deleted after the deletion time has expired.

5 **[0009B]** In another example aspect, there is provided a system comprising: one or more processors; and a non-transitory computer-readable medium containing instructions. When executed by the one or more processors, the instructions cause the one or more processors to perform operations including: receive, by a server, a reference data set associated with a media segment, wherein the server is configured to identify an unidentified media segment by matching an unidentified data set with the reference data set, and wherein the unidentified data set is associated with the unidentified media segment; store the reference data set in a reference database; receive popularity data indicating a popularity of the media segment, wherein the popularity data includes values associated with at least one or more of viewing information of the media segment, rating information of the media segment, or information associated with a posted message on a remote source, and wherein the posted message is associated with the media segment; determine a deletion time for the reference data set, wherein the deletion time is determined by calculating a representative value of two or more values of the popularity data, the representative value being one of an average, a mean, or a median of the two or more of the values of the popularity data; and delete the reference data set from the reference database, wherein the reference data set is deleted after the deletion time has expired.

10 **[0009C]** In another example aspect, there is provided a computer-program product tangibly embodied in a non-transitory machine-readable storage medium, including instructions that, when executed by the one or more processors, cause the one or more processors to: receive, by a server, a reference data set associated with a media segment, wherein the server is configured to identify an unidentified media segment by matching an unidentified data set with the reference data set, and wherein the unidentified data set is associated with the unidentified media segment; store the reference data set in a reference database; receive popularity data indicating a popularity of the media segment, wherein the popularity data includes values associated with at least one or more of viewing information of the media segment, rating information of the media segment, or information associated with a posted message on a remote source, and wherein the posted message is associated with the media segment; determine a deletion time for the reference data set, wherein the deletion time is determined by calculating a representative value of two or more values of the popularity data, the representative value being one of an average, a mean, or a median of the two or more of the values of the popularity data; and delete the reference data set

from the reference database, wherein the reference data set is deleted after the deletion time has expired.

5 [0010] The features, aspects, and advantages of the present disclosure will be most easily understood when the following description is read with reference to the accompanying drawings in which like numerical identifications represent like components or parts throughout the drawings.

BRIEF DESCRIPTION OF THE FIGURES

0 [0011] Illustrative examples are described in detail below with reference to the following figures.

[0012] FIG. 1 is a block diagram of an example of a matching system for identifying video content being viewed by a media system.

[0013] FIG. 2 illustrates an example of a matching system identifying unknown data points.

[0014] FIG. 3 is a block diagram of an example of a video capture system.

5 [0015] FIG. 4 is a block diagram of an example of a system for collecting video content presented by a display.

[0016] FIG. 5 illustrates an example of a matching system for optimizing resource utilization.

[0017] FIG. 6 is a diagram illustrating an example of relative weights for data used to determine whether to retain a reference data set.

10 [0018] FIG. 7 is a flowchart illustrating an example of a process for optimizing resource utilization in a matching system.

[0019] FIG. 8 is a chart illustrating point locations and the path points around them.

[0020] FIG. 9 is a chart illustrating a set of points that lie within distance from a query point.

[0021] FIG. 10 is a chart illustrating possible point values.

15 [0022] FIG. 11 is a chart illustrating a space divided into rings of exponentially growing width.

[0023] FIG. 12 is a chart illustrating self-intersecting paths and a query point.

[0024] FIG. 13 is a chart illustrating three consecutive point locations and the path points around them.

20

DETAILED DESCRIPTION

[0025] In the following description, for the purposes of explanation, specific details are set forth in order to provide a thorough understanding of examples of the invention. However, it will be apparent that various examples may be practiced without these specific details. The figures and description are not intended to be restrictive.

[0026] The ensuing description provides exemplary examples only, and is not intended to limit the scope, applicability, or configuration of the disclosure. Rather, the ensuing description of the exemplary examples will provide those skilled in the art with an enabling description for implementing an exemplary example. It should be understood that various changes may be made in the function and arrangement of elements without departing from the spirit and scope of the invention as set forth in the appended claims.

[0027] Specific details are given in the following description to provide a thorough understanding of the examples. However, it will be understood by one of ordinary skill in the art that the examples may be practiced without these specific details. For example, circuits, systems, networks, processes, and other components may be shown as components in block diagram form in order not to obscure the examples in unnecessary detail. In other instances, well-known circuits, processes, algorithms, structures, and techniques may be shown without unnecessary detail in order to avoid obscuring the examples.

[0028] Also, it is noted that individual examples may be described as a process which is depicted as a flowchart, a flow diagram, a data flow diagram, a structure diagram, or a block diagram. Although a flowchart may describe the operations as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be re-arranged. A process is terminated when its operations are completed, but could have additional steps not included in a figure. A process may correspond to a method, a function, a procedure, a subroutine, a subprogram, etc. When a process corresponds to a function, its termination can correspond to a return of the function to the calling function or the main function.

[0029] The term “machine-readable storage medium” or “computer-readable storage medium” includes, but is not limited to, portable or non-portable storage devices, optical storage devices, and various other mediums capable of storing, containing, or carrying instruction(s) and/or data. A machine-readable storage medium or computer-readable storage medium may include a non-transitory medium in which data can be stored and that does not include carrier waves and/or transitory electronic signals propagating wirelessly or over wired connections. Examples of a non-transitory medium may include, but are not limited to, a magnetic disk or tape, optical storage media such as compact disk (CD) or digital versatile disk (DVD), flash memory, memory

or memory devices. A computer-program product may include code and/or machine-executable instructions that may represent a procedure, a function, a subprogram, a program, a routine, a subroutine, a module, a software package, a class, or any combination of instructions, data structures, or program statements. A code segment may be coupled to another code segment or a hardware circuit by passing and/or receiving information, data, arguments, parameters, or memory contents. Information, arguments, parameters, data, or other information may be passed, forwarded, or transmitted using any suitable means including memory sharing, message passing, token passing, network transmission, or other transmission technique.

[0030] Furthermore, examples may be implemented by hardware, software, firmware, middleware, microcode, hardware description languages, or any combination thereof. When implemented in software, firmware, middleware or microcode, the program code or code segments to perform the necessary tasks (e.g., a computer-program product) may be stored in a machine-readable medium. A processor(s) may perform the necessary tasks.

[0031] Systems depicted in some of the figures may be provided in various configurations. In some examples, the systems may be configured as a distributed system where one or more components of the system are distributed across one or more networks in a cloud computing system.

[0032] As described in further detail below, certain aspects and features of the present disclosure relate to identifying unknown video segments by comparing unknown data points to one or more reference data points. The systems and methods described herein improve the efficiency of storing and searching large data sets that are used to identify the unknown video segments. For example, the systems and methods allow identification of the unknown data segments while reducing the density of the large data sets required to perform the identification. The techniques can be applied to any system that harvests and manipulates large volumes of data. Illustrative examples of these systems include automated content-based searching systems (e.g., automated content recognition for video-related applications or other suitable application), MapReduce systems, Bigtable systems, pattern recognition systems, facial recognition systems, classification systems, computer vision systems, data compression systems, cluster analysis, or any other suitable system. One of ordinary skill in the art will appreciate that the techniques described herein can be applied to any other system that stores data that is compared to unknown

data. In the context of automated content recognition (ACR), for example, the systems and methods reduce the amount of data that must be stored in order for a matching system to search and find relationships between unknown and known data groups.

5 [0033] By way of example only and without limitation, some examples described herein use an automated audio and/or video content recognition system for illustrative purposes. However, one of ordinary skill in the art will appreciate that the other systems can use the same techniques.

10 [0034] A significant challenge with ACR systems and other systems that use large volumes of data can be managing the amount of data that is required for the system to function. Another challenge includes a need to build and maintain a database of known content to serve as a reference to match incoming content. Building and maintaining such a database involves collecting and digesting a vast amount (e.g., hundreds, thousands, or more) of content (e.g., nationally distributed television programs and an even larger amount of local television broadcasts among many other potential content sources). The digesting can be performed using any available technique that reduces the raw data (e.g., video or audio) into compressed, 15 searchable data. With a 24-hour, seven-day-a-week operating schedule and a sliding window of perhaps two weeks of content (e.g., television programming) to store, the data volume required to perform ACR can build rapidly. Similar challenges can be present with other systems that harvest and manipulate large volumes of data, such as the example systems described above.

20 [0035] The systems and methods described herein can allow improved management of system resources in an ACR system. For example, examples can increase the efficiency of system resource utilization by managing the duration that data related to video segments are retained based on data that takes into account an identified popularity of a video segment. The identified popularity can be determined by algorithms that take into account numbers of viewers who watched the video segment, ratings of the video segment, metrics derived from social media, or 25 any other factor that can indicate likelihood that the video segment will be viewed. While description herein can refer to a video segment, other media segments can also be used, including audio segments.

30 [0036] FIG. 1 illustrates a matching system 100 that can identify unknown content. In some examples, the unknown content can include one or more unknown data points. In such examples, the matching system 100 can match unknown data points with reference data points to identify

unknown video segments associated with the unknown data points. The reference data points can be included in a reference database 116.

[0037] The matching system 100 includes a client device 102 and a matching server 104. The client device 102 includes a media client 106, an input device 108, an output device 110, and one or more contextual applications 126. The media client 106 (which can include a television system, a computer system, or other electronic device capable of connecting to the Internet) can decode data (e.g., broadcast signals, data packets, or other frame data) associated with video programs 128. The media client 106 can place the decoded contents of each frame of the video into a video frame buffer in preparation for display or for further processing of pixel information of the video frames. The client device 102 can be any electronic decoding system that can receive and decode a video signal. The client device 102 can receive video programs 128 and store video information in a video buffer (not shown). The client device 102 can process the video buffer information and produce unknown data points (which can referred to as “cues”), described in more detail below with respect to FIG. 3. The media client 106 can transmit the unknown data points to the matching server 104 for comparison with reference data points in the reference database 116.

[0038] The input device 108 can include any suitable device that allows a request or other information to be input to the media client 106. For example, the input device 108 can include a keyboard, a mouse, a voice-recognition input device, a wireless interface for receiving wireless input from a wireless device (e.g., from a remote controller, a mobile device, or other suitable wireless device), or any other suitable input device. The output device 110 can include any suitable device that can present or otherwise output information, such as a display, a wireless interface for transmitting a wireless output to a wireless device (e.g., to a mobile device or other suitable wireless device), a printer, or other suitable output device.

[0039] The matching system 100 can begin a process of identifying a video segment by first collecting data samples from known video data sources 118. For example, the matching server 104 collects data to build and maintain a reference database 116 from a variety of video data sources 118. The video data sources 118 can include media providers of television programs, movies, or any other suitable video source. Video data from the video data sources 118 can be provided as over-the-air broadcasts, as cable TV channels, as streaming sources from the

Internet, and from any other video data source. In some examples, the matching server 104 can process the received video from the video data sources 118 to generate and collect reference video data points in the reference database 116, as described below. In some examples, video programs from video data sources 118 can be processed by a reference video program ingest system (not shown), which can produce the reference video data points and send them to the reference database 116 for storage. The reference data points can be used as described above to determine information that is then used to analyze unknown data points.

[0040] The matching server 104 can store reference video data points for each video program received for a period of time (e.g., a number of days, a number of weeks, a number of months, or any other suitable period of time) in the reference database 116. The matching server 104 can build and continuously or periodically update the reference database 116 of television programming samples (e.g., including reference data points, which may also be referred to as cues or cue values). In some examples, the data collected is a compressed representation of the video information sampled from periodic video frames (e.g., every fifth video frame, every tenth video frame, every fifteenth video frame, or other suitable number of frames). In some examples, a number of bytes of data per frame (e.g., 25 bytes, 50 bytes, 75 bytes, 100 bytes, or any other amount of bytes per frame) are collected for each program source. Any number of program sources can be used to obtain video, such as 25 channels, 50 channels, 75 channels, 100 channels, 200 channels, or any other number of program sources. Using the example amount of data, the total data collected during a 24-hour period over three days becomes very large. Therefore, reducing the number of actual reference data point sets is advantageous in reducing the storage load of the matching server 104.

[0041] The media client 106 can send a communication 122 to a matching engine 112 of the matching server 104. The communication 122 can include a request for the matching engine 112 to identify unknown content. For example, the unknown content can include one or more unknown data points and the reference database 116 can include a plurality of reference data points. The matching engine 112 can identify the unknown content by matching the unknown data points to reference data in the reference database 116. In some examples, the unknown content can include unknown video data being presented by a display (for video-based ACR), a search query (for a MapReduce system, a Bigtable system, or other data storage system), an

unknown image of a face (for facial recognition), an unknown image of a pattern (for pattern recognition), or any other unknown data that can be matched against a database of reference data. The reference data points can be derived from data received from the video data sources 118. For example, data points can be extracted from the information provided from the video data sources 118 and can be indexed and stored in the reference database 116.

[0042] The matching engine 112 can send a request to the candidate determination engine 114 to determine candidate data points from the reference database 116. A candidate data point can be a reference data point that is a certain determined distance from the unknown data point. In some examples, a distance between a reference data point and an unknown data point can be determined by comparing one or more pixels (e.g., a single pixel, a value representing group of pixels (e.g., a mean, an average, a median, or other value), or other suitable number of pixels) of the reference data point with one or more pixels of the unknown data point. In some examples, a reference data point can be the certain determined distance from an unknown data point when the pixels at each sample location are within a particular pixel value range.

[0043] In one illustrative example, a pixel value of a pixel can include a red value, a green value, and a blue value (in a red-green-blue (RGB) color space). In such an example, a first pixel (or value representing a first group of pixels) can be compared to a second pixel (or value representing a second group of pixels) by comparing the corresponding red values, green values, and blue values respectively, and ensuring that the values are within a certain value range (e.g., within 0-5 values). For example, the first pixel can be matched with the second pixel when (1) a red value of the first pixel is within 5 values in a 0-255 value range (plus or minus) of a red value of the second pixel, (2) a green value of the first pixel is within 5 values in a 0-255 value range (plus or minus) of a green value of the second pixel, and (3) a blue value of the first pixel is within 5 values in a 0-255 value range (plus or minus) of a blue value of the second pixel. In such an example, a candidate data point is a reference data point that is an approximate match to the unknown data point, leading to multiple candidate data points (related to different media segments) being identified for the unknown data point. The candidate determination engine 114 can return the candidate data points to the matching engine 112.

[0044] For a candidate data point, the matching engine 112 can add a token into a bin that is associated with the candidate data point and that is assigned to an identified video segment from

which the candidate data point is derived. A corresponding token can be added to all bins that correspond to identified candidate data points. As more unknown data points (corresponding to the unknown content being viewed) are received by the matching server 104 from the client device 102, a similar candidate data point determination process can be performed, and tokens can be added to the bins corresponding to identified candidate data points. Only one of the bins corresponds to the segment of the unknown video content being viewed, with the other bins corresponding to candidate data points that are matched due to similar data point values (e.g., having similar pixel color values), but that do not correspond to the actual segment being viewed. The bin for the unknown video content segment being viewed will have more tokens assigned to it than other bins for segments that are not being watched. For example, as more unknown data points are received, a larger number of reference data points that correspond to the bin are identified as candidate data points, leading to more tokens being added to the bin. Once a bin includes a particular number of tokens, the matching engine 112 can determine that the video segment associated with the bin is currently being displayed on the client device 102. A video segment can include an entire video program or a portion of the video program. For example, a video segment can be a video program, a scene of a video program, one or more frames of a video program, or any other portion of a video program.

[0045] FIG. 2 illustrates components of a matching system 200 for identifying unknown data. For example, the matching engine 212 can perform a matching process for identifying unknown content (e.g., unknown media segments, a search query, an image of a face or a pattern, or the like) using a database of known content (e.g., known media segments, information stored in a database for searching against, known faces or patterns, or the like). For example, the matching engine 212 receives unknown data content 202 (which can be referred to as a “cue”) to be matched with a reference data point of the reference data points 204 in a reference database. The unknown data content 202 can also be received by the candidate determination engine 214, or sent to the candidate determination engine 214 from the matching engine 212. The candidate determination engine 214 can conduct a search process to identify candidate data points 206 by searching the reference data points 204 in the reference database. In one example, the search process can include a nearest neighbor search process to produce a set of neighboring values (that are a certain distance from the unknown values of the unknown data content 202). The candidate data points 206 are input to the matching engine 212 for conducting the matching

process to generate a matching result 208. Depending on the application, the matching result 208 can include video data being presented by a display, a search result, a determined face using facial recognition, a determined pattern using pattern recognition, or any other result.

[0046] In determining candidate data points 206 for an unknown data point (e.g., unknown data content 202), the candidate determination engine 214 determines a distance between the unknown data point and the reference data points 204 in the reference database. The reference data points that are a certain distance from the unknown data point are identified as the candidate data points 206. In some examples, a distance between a reference data point and an unknown data point can be determined by comparing one or more pixels of the reference data point with one or more pixels of the unknown data point, as described above with respect to FIG. 1. In some examples, a reference data point can be the certain distance from an unknown data point when the pixels at each sample location are within a particular value range. As described above, a candidate data point is a reference data point that is an approximate match to the unknown data point, and because of the approximate matching, multiple candidate data points (related to different media segments) are identified for the unknown data point. The candidate determination engine 114 can return the candidate data points to the matching engine 112.

[0047] FIG. 3 illustrates an example of a video ingest capture system 400 including a memory buffer 302 of a decoder. The decoder can be part of the matching server 104 or the media client 106. The decoder may not operate with or require a physical television display panel or device. The decoder can decode and, when required, decrypt a digital video program into an uncompressed bitmap representation of a television program. For purposes of building a reference database of reference video data (e.g., reference database 316), the matching server 104 can acquire one or more arrays of video pixels, which are read from the video frame buffer. An array of video pixels is referred to as a video patch. A video patch can be any arbitrary shape or pattern but, for the purposes of this specific example, is described as a 10x10 pixel array, including ten pixels horizontally by ten pixels vertically. Also for the purpose of this example, it is assumed that there are 25 pixel-patch positions extracted from within the video frame buffer that are evenly distributed within the boundaries of the buffer.

[0048] An example allocation of pixel patches (e.g., pixel patch 304) is shown in FIG. 3. As noted above, a pixel patch can include an array of pixels, such as a 10x10 array. For example, the

pixel patch 304 includes a 10x10 array of pixels. A pixel can include color values, such as a red, a green, and a blue value. For example, a pixel 306 is shown having Red-Green-Blue (RGB) color values. The color values for a pixel can be represented by an eight-bit binary value for each color. Other suitable color values that can be used to represent colors of a pixel include luma and chroma (Y, Cb, Cr) values or any other suitable color values.

[0049] A mean value (or an average value in some cases) of each pixel patch is taken, and a resulting data record is created and tagged with a time code (or time stamp). For example, a mean value is found for each 10x10 pixel patch array, in which case twenty-four bits of data per twenty-five display buffer locations are produced for a total of 600 bits of pixel information per frame. In one example, a mean of the pixel patch 304 is calculated, and is shown by pixel patch mean 308. In one illustrative example, the time code can include an “epoch time,” which representing the total elapsed time (in fractions of a second) since midnight, January 1, 1970. For example, the pixel patch mean 308 values are assembled with a time code 412. Epoch time is an accepted convention in computing systems, including, for example, Unix-based systems. Information about the video program, known as metadata, is appended to the data record. The metadata can include any information about a program, such as a program identifier, a program time, a program length, or any other information. The data record including the mean value of a pixel patch, the time code, and metadata, forms a “data point” (also referred to as a “cue”). The data point 310 is one example of a reference video data point.

[0050] A process of identifying unknown video segments begins with steps similar to creating the reference database. For example, FIG. 4 illustrates a video ingest capture system 400 including a memory buffer 402 of a decoder. The video ingest capture system 400 can be part of the client device 102 that processes data presented by a display (e.g., on an Internet-connected television monitor, such as a smart TV, a mobile device, or other television viewing device). The video ingest capture system 400 can utilize a similar process to generate unknown video data point 410 as that used by system 300 for creating reference video data point 310. In one example, the media client 106 can transmit the unknown video data point 410 to the matching engine 112 to identify a video segment associated with the unknown video data point 410 by the matching server 104.

[0051] As shown in FIG. 4, a video patch 404 can include a 10x10 array of pixels. The video patch 404 can be extracted from a video frame being presented by a display. A plurality of such pixel patches can be extracted from the video frame. In one illustrative example, if twenty-five such pixel patches are extracted from the video frame, the result will be a point representing a position in a 75-dimension space. A mean (or average) value can be computed for each color value of the array (e.g., RGB color value, Y, Cr, Cb color values, or the like). A data record (e.g., unknown video data point 410) is formed from the mean pixel values and the current time is appended to the data. One or more unknown video data points can be sent to the matching server 104 to be matched with data from the reference database 116 using the techniques described above.

[0052] FIG. 5 illustrates an example of a matching system 500 for optimizing resource utilization. The matching system 500 can be similar to the matching system 100 described in FIG. 1. For example, the matching system 500 can include a client device 510 (which can be the same or similar to the client device 102). The client device 510 can receive frames of a video segment from a remote source. In some examples, the frames can be in the form of a broadcast signal. In other examples, the frames can be in the form of data (e.g., frame data) sent to the client device 510 in packets using a network (e.g., the Internet). In some examples, a frame can include pixel values that are used by the client device 510 to present the frame to a user on a display.

[0053] In some examples, the pixel values can be sent directly to a matching server 520 (which can perform similar to the matching server 104) from the client device 510. In other examples, a portion of the pixel values can be sent to the matching server 520 using techniques described above (e.g., pixel cue points can be sent to the matching server 520). In other examples, the pixel values (or a portion thereof) can be summarized by the client device 510, where the summarized pixel values can be sent to the matching server 520.

[0054] The matching server 520 can include a matching engine 530 (which can perform the same or similar to the matching engine 112) to identify an unidentified video segment being displayed on the client device 510. To identify the unidentified video segment, the matching engine 530 can receive pixel values (or pixel cue points, average pixel values, or other representations of a frame) associated with one or more frames from the client device 510. The

matching engine 530 can also receive or obtain reference data to match with the pixel values. In some examples, the reference data can be received from a reference database 550. In such examples, the reference database 550 can send one or more reference data sets, which can be stored in the reference database 550, to the matching engine 530. In other examples, the matching engine 530 can obtain the one or more reference data sets from the reference database 550 using a database query. In some examples, the matching engine 530 can send a request to a candidate determination engine 540 to determine candidate data points from the reference database 550. As described above, the candidate data points can be reference data points that are a certain determined distance from one or more data points in a frame. The candidate determination engine 540 can return the candidate data points to the matching engine 530.

[0055] As described above, for each of the candidate data points, the matching engine 530 can add a token into a bin corresponding to an identified video segment. The identified video segment can be associated with one or more candidate data points that include the candidate data point causing the token to be added to the bin. This process can repeat for each new frame that is sent to the matching engine 530 from the client device 510. Once a bin includes a particular number of tokens, the matching engine 530 can determine that the video segment associated with the bin is currently being displayed by the client device 510.

[0056] Once the matching engine 530 identifies a video segment, the matching engine 530 can send a signal to a results engine 580 (which can perform similar to the results engine 120) indicating the identified video segment. The results engine 580 can determine an action to perform in response to the identified video segment. In some examples, the results engine 580 can identify a contextual application that corresponds to the identified video segment. The contextual application can be stored on the client device 510. In some examples, the contextual application can be an application associated with at least a portion of a video segment. The contextual application can run on the client device 510 to display content on the client device 510. The content displayed can be an overlay on content already being displayed on the client device 510 or can replace the content being displayed on the client device 510.

[0057] The results engine 580 can send a signal to the client device 510 to cause the identified contextual application to perform an action. For example, the contextual application can display content on the client device 510. The content can be related to the identified video segment.

Other actions the identified contextual application can perform include, but are not limited to, replacement of a commercial message with one more commercial messages directed to a specific viewer, additional information regarding the identified video segment, or an opportunity to interact with the identified video segment or other's watching the identified video segment.

5 [0058] In some examples, the reference data sets in the reference database 550 can be managed by a retention engine 560. Management of the reference data sets can include determining when to remove, or delete, reference data sets from the reference database 550. In some examples, when a reference data set that is associated with a video segment is removed, one or more other reference data sets that are also associated with the video segment can be removed. In some
10 examples, the one or more other reference data sets are not removed. Further details regarding removal or deletion of reference data sets are described below.

[0059] In some examples, reference data sets that are removed, or deleted, from the reference database 550 can be added to a backup database (not shown). In some examples, the backup database can be located on the matching server 520. In other examples, the backup database can
15 be located remotely from the matching server 520. In such examples, the backup database can be on a cloud (e.g., one or more remote servers associated with the matching server 520). The candidate determination engine 540 can be configured to search the backup database for a matching candidate data set after the reference database 550 has been searched and no match is found or an insufficient number of matches has been found.

20 [0060] A person of ordinary skill will recognize that there might be more than one backup database. Zero or more of the backup databases can be located with the matching server 520 and zero or more the backup databases can be located remotely from the matching server 520. In some examples, each additional backup database can retain reference data sets longer than a previous backup database. When a reference data set is removed completely from databases
25 associated with the matching server 520, the reference data set would no longer be able to be compared to a unidentified data set by the candidate determination engine 540. In some examples, each backup database can be managed by the retention engine 560. In other examples, each backup database can be managed by a separate retention engine. A backup database can either use the same length of time as the reference database (essentially doubling the life of a

reference data set) or can use a different length of time based on a same, a similar, or a different process as that used to determine the length of time for the reference database.

[0061] A person of ordinary skill in the art will recognize that a number of backup databases used will depend on resources available for an enterprise managing the backup databases, as well as efficiency concerns with having the candidate determining engine searching more and more backup databases. In some examples, each backup database can be searched using a separate process, allowing searches of multiple backup databases to minimally affect normal operation of the candidate determination engine 540. In some examples, the separate processes can be run in parallel so that that multiple backup databases can be searched at least partially in parallel. For example, a first database can be searched until a certain percent of the first database has been searched, at which time a second database can begin being searched in parallel with the first database.

[0062] The retention engine 560 can determine a length of time to retain reference data sets in the reference database 550 (e.g., referred to as a deletion time, a retention time, a removal time, or any other term that indicates a time at which a reference data set is removed, or deleted, from the reference database 550). The retention engine 560 can use popularity data (e.g., statistics, numbers, information, or other popularity data) associated with a video segment to determine the length of time to retain a reference data set.

[0063] In some examples, the length of time can be different for reference data sets associated with different video segments, for each reference data set, or any combination thereof. A person of ordinary skill in the art will recognize that many configurations of length of time for reference data sets can apply.

[0064] In some examples, the length of time can be extended whenever a particular number of views of a video segment occur with the matching server 520 (e.g., 1 view per day, 10 views per day, 1 view per week, an average of 10 views per week, or any other suitable number of views). The views of the video segment can be determined when the matching server 520 identifies that a media system is displaying the video segment, as will be discussed below. In some examples the length of time can be reassessed when the length of time expires.

[0065] In some examples, the length of time can be consistent for all video segments of a content item (e.g., a movie, a show, an episode, etc.). For example, all reference data sets for all video segments of a content item can be retained for a day, a week, or another amount of time. In some examples, the length of time can be consistent for all reference data sets associated with a single video segment. In such examples, even though a reference data set is not matched by the matching server 520, the reference data set can be retained as long as another reference data set associated with the same video segment is matched by the matching server 520.

[0066] In some examples, the popularity data can be associated with an identified popularity of a video segment. The popularity data can include viewing information, which can indicate a number of views of a video segment from a viewing information source 572. In some examples, the number of views can be associated with a particular time period. In some examples, the time period can be the day that a video segment is broadcast (e.g., same-day viewing data) or a period of time later after the day that a video segment is broadcast (delayed viewing data). A person of ordinary skill in the art will recognize that other time periods can be used.

[0067] In some examples, the viewing information source 572 can be a third-party server (e.g., over a network, such as using the Internet 570). Examples of viewing information can include any type of audience measurement system, including, but not limited to, Nielsen/NetRatings, Nielsen BuzzMetrics, Nielsen Media Research (NMR) comScore, Wakoopa, and Hitwise, Visible Measures, GfK, Sightcorp, TruMedia, Quividi, reLEYEble, stickyPiXEL, Cognitec, goCount and CognoVision, or any other form of audience measurement system.

[0068] In some examples, the viewing information can be determined using the matching engine 530. For example, when a video segment is identified by the matching engine 530 (when a bin reaches a certain number of tokens, as described above), the matching server 520 can store a record that the video segment has been displayed by a media system. In such an example, the matching server 520 can keep track of the number of times a video segment has been displayed by a media system. Thus, the number of times a video segment has been displayed by a media system can be used as a number of viewings of a video segment for purposes of the viewing information.

[0069] A person of ordinary skill in the art will also recognize that more than one viewing information can be used to determine the length of time. For example, viewing information using

the matching engine 530 and viewing information from a viewing information source 572 can be used. For another example, viewing information from a first viewing information source and viewing information from a second viewing information source can be used.

5 [0070] The popularity data can also include rating information of the video segment from a rating source 574. The rating information can indicate a desirability, popularity, and/or demand of the video segment. For example, the rating information can be a percentage of viewers who liked the video segment. The percentage of viewers who liked the video segment can be determined by a source on a network (such as the Internet 570) that provides a rating for video segments. For example, the rating can be a rating given to the video segment by one or more
10 reviewers on a website, a mobile application, or the like. In some examples, the one or more reviewers can be from at least one or more of a professional group of reviewers and a recreational group of reviewers. The rating sources 574 can be a third-party server (e.g., over a network, such as using the Internet 570). Examples of the rating source 574 can include Rotten Tomatoes, Metacritic, Everyone's a Critic, Reviewer, Movie Attractions, Flixster, FilmCrave,
15 Flickchart, blogs, and systems such as the "undulating curve of shifting expectations" (which use automated systems to review video segments). A person of ordinary skill in the art will also recognize that more than one rating information or source can be used to determine the length of time. In one illustrative example, rating information using a professional reviewer and rating information from a recreational group of reviewers can be used.

20 [0071] The popularity data can also include information associated with one or more posts (e.g., a posted message) on one or more remote sources (e.g., a review source or a social media source 576, such as a website and a mobile application). The one or more posts can be associated with the video segment. The social media source 576 can be a third-party server (e.g., over a network, such as using the Internet 570). Examples of a social media source can include, but are
25 not limited to, Facebook, YouTube, Twitter, LinkedIn, Pinterest, Google Plus+, Instagram, Reddit, and Rotten Tomatoes.

[0072] Examples of information associated with one or more posts can include, but are not limited to, at least one or more of a number of posted messages from one or more remote sources, a number of posted messages from one or more remote sources that are positive in
30 nature toward the video segment (based on a natural language parsing of the one or more

message), a number of posted messages from one or more remote sources that are negative in nature toward the video segment (again, based on a natural language parsing of the one or more message), a number of positive indications from one or more remote sources (e.g., a “Like” on Facebook), a number of reposts of a posted message associated with a video segment on one or more remote sources, a number of views of a page associated with the video segment on one or more remote sources, or any combination thereof. In some examples, a rating (e.g., 1 to 100) can be determined from one or more posted messages that identifies how much the video segment is liked based on words used in the one or more posted messages. In such examples, certain words can score certain points. In some examples, a number of mentions on a social media platform regarding a video segment or an actor associated with the video program.

[0073] A person of ordinary skill in the art will also recognize that more than one information associated with one or more posts on one or more remote sources information can be used to determine the length of time. For example, information of a number of posted messages from one or more remote sources and a number of positive indications from the one or more remote sources can be used.

[0074] In some examples, the length of time can be determined by calculating a representative value (e.g., an average, a mean, a median, or some combination thereof) of the popularity data. The representative value can indicate a probability of whether a reference data set will be (or is sufficiently likely to be) needed by the matching system in the future to justify a cost of retaining the reference data set in the reference database 550. In some examples, an average can indicate a likelihood that a video segment will be displayed on a media system (either associated with the matching server or not) within a particular amount of time (e.g., within a day, within a week, within a month, or with any other suitable period of time). When calculating the average, the retention engine 560 can identify the popularity data available, and use the available popularity data. While calculating an average, if a video segment is not included in a source, a reference data set associated with the video segment can receive a zero for the source.

[0075] In some examples, the popularity data can be normalized such that the popularity data is all based on a similar or a same scale. For example, a rating from 0 to 1 can be multiplied by 100 to correspond to a rating from 0 to 100. For another example, a largest number for each type of popularity data can either be determined by analyzing a type of popularity data or pre-

identified by a user. In such an example, a number associated with a type of popularity data can be divided by the largest number such that all of the numbers are a percentage of the largest number. In some examples, numbers larger than the largest number can either be allowed to pass a percentage of 1, or the number can be capped at 1.

5 [0076] In some examples, weights can be applied to the popularity data, such that some portions of the popularity data affect the average more than others. In an illustrative example, weighting of popularity data can include: direct mention of a video segment or actor associated with the video segment could be weighted higher than a mention of the video segment on a social media communication. Another weighting factor could be a number of words used in a
10 message associated with a video segment. Another weighting factor could be a number of specific actors associated with a video segment cited in a message. FIG. 6 is a diagram illustrating an example of relative weights for popularity data used to determine whether to retain a reference data set. A person of ordinary skill in the art will recognize that a weight can be within a range, such as a range from 0 to 100, a range from 0 to 1, or other suitable range of
15 weight values. The sum of all of the weights can equal to 100, to 1, or other value at the top of a range. The weights can be multiplied by portions of the popularity data to either enhance the popularity data or reduce the popularity data. Other methods of weighting data can also be used. While FIG. 6 illustrates a particular order of relative weights, a person of ordinary skill in the art will recognize that a different order can be used and different popularity data can be used.

20 [0077] In one example, a highest weight (e.g., 50%) can be associated with one or more popularity datum of the popularity data. A popularity datum can be a viewing information from the matching server (block 620). The viewing information can indicate a number of times a video segment is identified by the matching server as having been viewed by a media system associated with the matching system. The number can increase every time that the matching
25 server identifies the video segment as being displayed by a media system, as described above.

[0078] In the same example, a medium weight (e.g., 30%) can be associated with one or more popularity datum of the popularity data. For example, a first popularity datum of the popularity data can be a list of the top 20 programs for same-day viewing information (block 630). The list can include a number of times a top 20 program is viewed. The same-day viewing information
30 can be received from a secondary source (e.g., from a viewing information source 572 over the

Internet). The same-day viewing information can be weighted less than the viewing information from the matching server (block 620) because the same-day viewing information is not from the matching server, but from a secondary source. A second popularity datum of the popularity data can be social media popularity information (block 650). The social media popularity information can be a number of posted messages regarding the video segment on one or more social media sources. For example, a wall post on Facebook can mention a video segment by name. The wall post can count as one posted message regarding the video segment.

[0079] In the same example, a low weight (e.g., 20%) can also be associated with one or more popularity datum of the popularity data. For example, a first popularity datum of the popularity data can be a second social media popularity information (block 640). The second social media popularity information can be a number of positive indications from one or more social media sources. A second popularity data of the popularity data can be a number of positive indications from the one or more remote sources (block 660). A person of ordinary skill will recognize that other configurations of weight can be used. In addition, any combination of popularity data can be used for a calculation of the weights.

[0080] The weights can be applied to the popularity data or the averages of the popularity data to create a weighted average of popularity data for a video segment (block 670). For example, weights can be applied by multiplication with an average data associated with a weight.

[0081] A total number (e.g., an average, a weighted average, or a sum of the popularity data data) associated with a video segment can be used to determine a length of time to retain one or more reference data sets associated with the video segment. In some examples, the length of time can be computed using a range for the total number. In some examples, a range can be predetermined (e.g., 0 to 50 or some other range that is received by the retention engine), based on particular percentages of the maximum total number possible, based on total numbers that have been seen by the matching system, or some other method of determining a range. In one illustrative example of percentages of the maximum total number possible, a maximum number for each popularity datum used to compute the total number can be summed together to determine the maximum total number possible. Then, each range can represent a certain percentage of the maximum total number possible. For example, a first range can be all total numbers that fall between 0 and 50% of the maximum total number possible.

[0082] In one illustrative example, when the total number is within a first range, a reference data set associated with the video segment associated with the total number can be retained for a first period of time (e.g., a day). When the total number is within a second range, the reference data set associated with the video segment associated with the total number can be retained for a second period of time (e.g., one week). A person of ordinary skill in the art will recognize that the ranges, the periods of time, and the number of ranges can be different than illustrated above.

[0083] The process of determining lengths of time can be dynamically configured by the retention engine 560. For example, the ranges can be manipulated during normal processing to see whether different ranges provide an effective balance between number of matches and size of the reference database. Effectiveness of the length of time determination can be based on a number of matches determined and an average time between matches in the reference database (where the average time is a positive whole number such as nanoseconds). For example, an efficiency measure can be calculated by dividing the number of matches determined by the average time. The efficiency measure can be optimized by changing the ranges, the periods of times, and the number of ranges by small increments until an optimized efficiency measure is found. The small increments can be tested by rematching past popularity data with different retention policies (e.g., various configurations of ranges, periods of time, and number of ranges) to see if the efficiency measure increases or decreases.

[0084] In some examples, a length of time can be determined based on viewing statistics that measured an average number of days that a video segment was viewed. Such data can be commercially available from companies such as AC Nielsen and Rentrack. The data can then be correlated with the popularity data (perhaps weighted by genre, e.g., sports, reality TV, crime series, etc.). The correlated data can then provide a direct time metric relative to popularity.

[0085] In other examples, a length of time can be determined based on tracking value of a weighted average. For example, when said an average value falls below a predetermined limit (e.g., below 50%), the video segment data can cease to be retained for a predetermined period of time.

[0086] FIG. 7 is a flowchart illustrating an example of a process 700 for optimizing resource utilization in a matching system. Process 700 is illustrated as a logical flow diagram, the operation of which represent a sequence of operations that can be implemented in hardware,

computer instructions, or a combination thereof. In the context of computer instructions, the operations represent computer-executable instructions stored on one or more computer-readable storage media that, when executed by one or more processors, perform the recited operations. Generally, computer-executable instructions include routines, programs, objects, components, data structures, and the like that perform particular functions or implement particular data types. The order in which the operations are described is not intended to be construed as a limitation, and any number of the described operations can be combined in any order and/or in parallel to implement the processes.

[0087] Additionally, the process 700 can be performed under the control of one or more computer systems configured with executable instructions and can be implemented as code (e.g., executable instructions, one or more computer programs, or one or more applications) executing collectively on one or more processors, by hardware, or combinations thereof. As noted above, the code can be stored on a machine-readable storage medium, for example, in the form of a computer program comprising a plurality of instructions executable by one or more processors. The machine-readable storage medium can be non-transitory.

[0088] At step 710, the process 700 includes receiving a reference data set associated with a media segment. In some examples, the reference data set can be received by a server. The server can be configured to identify an unidentified media segment by matching an unidentified data set with the reference data set. In some examples, the unidentified data set can be associated with the unidentified media segment. In some examples, the reference data set can include pixel data or audio data for a frame of the media segment.

[0089] At step 720, the process 700 includes storing the reference data set in a reference database.

[0090] At step 730, the process 700 includes receiving popularity data indicating a popularity of the media segment. In some examples, the popularity data includes at least one or more of viewing information of the media segment, rating information of the media segment, or information associated with a posted message on a remote source. In such examples, the posted message can be associated with the media segment. In some examples, the information associated with the posted message on the remote source can include at least one or more of a number of posted messages on the remote source, a number of the posted messages that are

positive, a number of the posted messages that are negative, a number of positive indications to the posted message, a number of reposts of the posted message, or a number of views of the posted message. In some examples, the posted messages can be associated with the media segment. In some examples, the viewing information can include a number of times the media
5 segment has been identified by the server.

[0091] At step 740, the process 700 includes determining a deletion time for the reference data set. In some examples, the deletion time can be determined using the popularity data. In some examples, determining a deletion time can include computing an average of two or more values of the popularity data. In such examples, each value of the popularity data can be associated with
10 a weight.

[0092] At step 750, the process 700 includes deleting the reference data set from the reference database. In some examples, the reference data set can be deleted after the deletion time has expired.

[0093] In some implementations, the process 700 can further include adding the reference set to a backup database when the reference data set is deleted from the reference database. In such implementations, the server can be configured to search the backup database after the reference
15 database for a data set matching the unidentified data set.

[0094] As discussed above, a video matching system can be configured to identify a media content stream when the media content stream includes an unscheduled media segment. As
20 further discussed above, identifying the media content stream may include identifying media content played by a media display device before or after the unscheduled media segment. Processes for identifying media content are discussed above with respect to FIG. 1. Specifically, the video content system may use samples taken from the display mechanism of the media content device (e.g., graphic and/or audio samples) and generate cues from these samples. The
25 video matching system may then match the cues against a reference database, where the database contains cues of known content.

[0095] The video matching system may further include various methods to improve the efficiency of finding potential matches, or “candidates” in the database. The database may contain an enormous number of cues, and thus the video matching system may include

algorithms for finding candidate cues to match against cues generated from the media content device's display mechanism. Locating candidate cues may be more efficient than other methods for matching cue values against the values in the database, such as matching a cue against every entry in the database.

- 5 [0096] Nearest neighbor and path pursuit are examples of techniques that can be used to locate candidate queues in the reference database. Below, an example of tracking video transmission using ambiguous cues is given, but the general concept can be applied to any field where match candidates are to be selected from a reference database.

[0097] A method for efficient video pursuit is presented. Given a large number of video
10 segments, the system must be able to identify in real time what segment a given query video input is taken from and in what time offset. The segment and offset together are referred to as the location. The method is called video pursuit since it must be able to efficiently detect and adapt to pausing, fast forwarding, rewinding, abrupt switching to other segments and switching to unknown segments. Before being able to pursue live video the database is processed. Visual cues
15 (a handful of pixel values) are taken from frames every constant fraction of a second and put in specialized data structure (note that this can also be done in real time). The video pursuit is performed by continuously receiving cues from the input video and updating a set of beliefs or estimates about its current location. Each cue either agrees or disagrees with the estimates, and they are adjusted to reflect the new evidence. A video location is assumed to be the correct one if
20 the confidence in this being true is high enough. By tracking only a small set of possible "suspect" locations, this can be done efficiently.

[0098] A method is described for video pursuit but uses mathematical constructs to explain and investigate it. It is the aim of this introduction to give the reader the necessary tools to translate between the two domains. A video signal, or video data, is comprised of sequential frames. Each
25 can be thought of as a still image. Every frame is a raster of pixels. Each pixel is made out of three intensity values corresponding to the red, green, and blue (RGB) make of that pixel's color. In the terminology used herein, a cue is a list of RGB values of a subset of the pixels in a frame and a corresponding time stamp. The number of pixels in a cue is significantly smaller than in a frame, usually between 5 and 15. Being an ordered list of scalar values, the cue values are in fact
30 a vector. This vector is also referred to as a point.

[0099] Although these points are in high dimension, usually between 15 and 150, they can be imagined as points in two dimensions. In fact, the illustrations will be given as two dimensional plots. Now, consider the progression of a video and its corresponding cue points. Usually a small change in time produces a small change in pixel values. The pixel point can be viewed as “moving” a little between frames. Following these tiny movements from frame to frame, the cue follows a path in space like a bead would on a bent wire.

[0100] In the language of this analogy, in video pursuit the locations of the bead in space (the cue points) are received and the part of wire (path) the bead is following is looked for. This is made significantly harder by two facts. First, the bead does not follow the wire exactly but rather keeps some varying unknown distance from it. Second, the wires are all tangled together. These statements are made exact in section 2. The algorithm described below does this in two conceptual steps. When a cue is received, the algorithm looks for all points on all the known paths that are sufficiently close to the cue point; these are called suspects. This is done efficiently using the Probabilistic Point Location in Equal Balls algorithm. These suspects are added to a history data structure and the probability of each of them indicating the true location is calculated. This step also includes removing suspect locations that are sufficiently unlikely. This history update process ensures that on the one hand only a small history is kept but on the other hand no probable locations are ever deleted. The generic algorithm is given in Algorithm 1 and illustrated in FIG. 8.

Algorithm 1 Generic path pursuit algorithm.

```

1: Set of suspects is empty
2: loop
3:   Receive latest cue.
4:   Find path points who are close to it.
5:   Add them to the set of suspects.
6:   Based on the suspects update the location likelihood function.
7:   Remove from suspect set those who do not contribute to the
   likelihood function.
8:   if A location is significantly likely then
9:     Output the likely location.
10:  end if
11: end loop

```

20

[0101] The following sections begin with describing the Probabilistic Point Location in Equal Balls (PPLEB) algorithm in Section 1. The PPLEB algorithm is used in order to perform line 5 in

Algorithm 1 above efficiently. The ability to perform this search for suspects quickly is crucial for the applicability of this method. In section 2 one possible statistical model is described for performing lines 6 and 7. The described model is a natural choice for the setup. It is also shown how it can be used very efficiently.

5 **[0102]** Section 1 - Probabilistic Point Location in Equal Balls

[0103] The following section describes a simple algorithm for performing probabilistic point location in equal balls (PPLEB). In the traditional PLEB (point location in equal balls), one starts with a set of n points x_i in \mathbb{R}^d and a specified ball of radius r . The algorithm is given $O(\text{poly}(n))$ preprocessing time to produce an efficient data structure. Then, given a query point x the
 10 algorithm is required to return all points x_i such that $\|x - x_i\| \leq r$. The set of points such that $\|x - x_i\| \leq r$ geometrically lie within a ball of radius r surrounding the query x (see FIG. 23). This relation is referred to as x_i being close to x or as x_i and x being neighbors.

[0104] The problem of PPLEB and the problem of nearest neighbor search are two similar problems that received much attention in the academic community. In fact, these problems were
 15 among the first studied in the field of computational geometry. Many different methods cater to the case where the ambient dimension is small or constant. These partition the space in different ways and recursively search through the parts. These methods include KD-trees, cover-trees, and others. Although very efficient in low dimension, when the ambient dimension is high, they tend to perform very poorly. This is known as the "curse of dimensionality". Various approaches
 20 attempt to solve this problem while overcoming the curse of dimensionality. The algorithm used herein uses a simpler and faster version of the algorithm and can rely on Local Sensitive Hashing.

[0105] Section 1.1 Locality Sensitive Hashing

[0106] In the scheme of local sensitive hashing, one devises a family of hash functions H such
 25 that:

$$Pr_{u \sim U}(u(x) \neq u(y) \mid \|x - y\| \leq r) \equiv p$$

$$Pr_{u \sim U}(u(x) \neq u(y) \mid \|x - y\| \geq 2r) \geq 2p$$

[0107] In words, the probability of x and y being mapped to the same value by h is significantly higher if they are close to each other.

5 [0108] For the sake of clarity, let us first deal with a simplified scenario where all incoming vectors are of the same length r' and $r' > \sqrt{2r}$. The reason for the latter condition will become clear later. First a random function $u \in U$ is defined, which separates between x and y according to the angle between them. Let \vec{u} be a random vector chosen uniformly from the unit sphere S^{d-1} and let $u(x) = \text{sign}(\vec{u} \cdot x)$. It is easy to verify that $Pr_{u \sim U}(u(x) \neq u(y)) = \theta_{x,y}/\pi$. Moreover, for
 10 any points x, y, x', y' on a circle such that $\|x' - y'\| \leq 2\|x - y\|$, $\theta_{x',y'} \leq 2\theta_{x,y}$ is achieved. Defining p , the following equations are used:

$$Pr_{u \sim U}(u(x) \neq u(y) \mid \|x - y\| \leq r) \equiv p$$

$$Pr_{u \sim U}(u(x) \neq u(y) \mid \|x - y\| \geq 2r) \geq 2p$$

[0109] The family of functions H is set to be a cross product of t independent copies of u , i.e. $h(x) = [u_1(x), \dots, u_t(x)]$. Intuitively, one would like to have that if $h(x) = h(y)$ then x and y are
 15 likely to be close to each other. Let us quantify that. First, compute the expected number of false positive mistakes n_{fp} . These are the cases for which $h(x) = h(y)$ but $\|x - y\| > 2r$. A value t is found for which n_{fp} is no more than 1, i.e. one is not expected to be wrong.

$$E[n_{fp}] \leq n(1 - 2p)^t \leq 1$$

$$\rightarrow t \geq \log(1/n) / \log(1 - 2p)$$

[0110] Now, the probability that $h(x) = h(y)$ given that they are neighbors is computed:

$$\begin{aligned}
 Pr(h(x) = h(y) \mid \|x - y\| \leq r) &\geq (1 - p)^{\log(1/r)/\log(1-2p)} \\
 &= (1/n)^{\log(1-p)/\log(1-2p)} \\
 &\geq 1/\sqrt{n}.
 \end{aligned}$$

[0111] Note here that one must have that $2p < 1$ which requires $r' > \sqrt{2r}$. This might not sound like a very high success probability. Indeed, $1/\sqrt{n}$ is significantly smaller than $1/2$. The next section will describe how to boost this probability up to $1/2$.

5

[0112] Section 1.2 The Point Search Algorithm

[0113] Each function h maps every point in space to a bucket. Define the bucket function $B_h: \mathbb{R}^d \rightarrow 2^{[n]}$ of a point x with respect to hash function h as $B_h(x) \equiv \{x_i \mid h(x_i) = h(x)\}$. The data structure maintained is $m = O(\sqrt{n})$ instances of bucket functions $[B_{h_1}, \dots, B_{h_m}]$. When one searches for a point x , the function returns $B(x) = \cup_i B_{h_i}(x)$. According to the previous section, there are two desired results:

10

$$\begin{aligned}
 Pr(x_j \in B(x) \mid \|x_j - x\| \leq r) &\geq 1/2 \\
 E[|B(x) \cap \{x_j \mid \|x - x_j\| > 2r\}|] &\leq \sqrt{n}.
 \end{aligned}$$

[0114] In other words, while with probability at least $1/2$ each neighbor of x is found, one is not likely to find many non-neighbors.

15

[0115] Section 1.3 Dealing with Different Radii Input Vectors

[0116] The previous sections only dealt with searching through vectors of the same length, namely r' . Now described is how one can use the construction as a building block to support a search in different radii. As seen in FIG. 11, the space is divided into rings of exponentially growing width. Ring i , denoted by R_i , includes all points x_i such that $\|x_i\| \in [2r(1+\epsilon)^i, 2r(1+\epsilon)^{i+1})$.

$)^{i+1}$]. Doing this achieves two ends. First, if x_i and x_j belong to the same ring, then $\|x_j\|/(1+\epsilon) \leq \|x_i\| \leq \|x_j\|(1+\epsilon)$. Second, any search can be performed in at most $1/\epsilon$ such rings. Moreover, if the maximal length vector in the data set is r' then the total number of rings in the system is $O(\log(r'/r))$.

5 [0117] Section 2 The Path Pursuit Problem

[0118] In the path pursuit problem, a fixed path in space is given along with the positions of a particle in a sequence of time points. The terms particle, cue, and point will be used interchangeably. The algorithm is required to output the position of the particle on the path. This is made harder by a few factors: the particle only follows the path approximately; the path can be
10 discontinuous and intersect itself many times; both particle and path positions are given in a sequence of time points (different for each).

[0119] It is important to note that this problem can simulate tracking a particle on any number of paths. This is simply done by concatenating the paths into one long path and interpreting the resulting position as the position on the individual paths.

15 [0120] More precisely, let path P be parametric curve $P: \mathbb{R} \rightarrow \mathbb{R}^d$. The curve parameter will be referred to as the time. The points on the path that are known to us are given in arbitrary time points t_i , i.e. n pairs $(t_i, P(t_i))$ are given. The particle follows the path but its positions are given in different time points, as shown in FIG. 12. Further, m pairs $(t'_j, x(t'_j))$ are given, where $x(t'_j)$ is the position of the particle in time t'_j .

20 [0121] Section 2.1 Likelihood Estimation

[0122] Since the particle does not follow the path exactly and since the path can intersect itself many times it is usually impossible to positively identify the position on the path the particle is actually on. Therefore, a probability distribution is computed on all possible path locations. If a location probability is significantly probable, the particle position is assumed to be known. The
25 following section describes how this can be done efficiently.

[0123] If the particle is following the path, then the time difference between the particle time stamp and the offset of the corresponding points on P should be relatively fixed. In other words,

if $x(t')$ is currently in offset t on the path then it should be close to $P(t)$. Also, τ seconds ago it should have been in offset $t - \tau$. Thus $x(t' - \tau)$ should be close to $P(t - \tau)$ (note that if the particle is intersecting the path, and $x(t')$ is close to $P(t)$ temporarily, it is unlikely that $x(t' - \tau)$ and $P(t - \tau)$ will also be close). Define the relative offset as $\Delta = t - t'$. Notice that as long as the particle is following the path the relative offset Δ remains unchanged. Namely, $x(t')$ is close to $P(t' + \Delta)$.

[0124] The maximum likelihood relative offset is obtained by calculating:

$$\Delta = \underset{\delta}{\operatorname{argmax}} Pr(x(t'_m), x(t'_{m-1}), \dots, x(t'_1) | P, \delta)$$

[0125] In words, the most likely relative offset is the one for which the history of the particle is most likely. This equation however cannot be solved without a statistical model. This model must quantify: how tightly x follows the path; how likely it is that x jumps between locations; and how smooth the path and particle curves are between the measured points.

[0126] Section 2.2 Time Discounted Binning

[0127] Now described is a statistical model for estimating the likelihood function. The model makes the assumption that the particle's deviation away from the path distributes normally with standard deviation σ . It also assumes that at any given point in time, there is some non-zero probability the particle will abruptly switch to another path. This is manifested by an exponential discount with time for past points. Apart from being a reasonable choice for a modeling point of view this model also has the advantage of being efficiently updateable. For some constant time unit 1: set the likelihood function to be proportional to f which is defined as follows:

$$f_m([\delta/\tau]) = \sum_{j=1}^m \sum_{i=1}^n e^{-\left(\frac{[x(t'_j) - P(t_i + \delta)]}{\sigma r}\right)^2} (1 - \zeta)^{t'_j - t'_i}$$

[0128] Here $\alpha \ll 1$ is a scale coefficient and $\zeta > 0$ is the probability that the particle will jump to a random location on the path in a given time unit.

[0129] Updating the function f efficiently can be achieved using the following simple observation.

$$f_m([\delta/\tau]) = \sum_{i=1}^n e^{-\left(\frac{\|x(t'_m) - P(t_i + \delta)\|}{\sigma r}\right)^2} + f_{m-1}([\delta/\tau])(1 - \xi)^{t'_m - t'_{m-1}}$$

5

[0130] Moreover, since $\alpha \ll 1$, if $\|x(t'_m) - P(t_i)\| \geq r$, the following occurs:

$$e^{-\left(\frac{\|x(t'_m) - P(t_i)\|}{\sigma r}\right)^2} \approx 0.$$

[0131] This is an important property of the likelihood function since the sum update can now
 10 performed only over the neighbors of $x(t'_i)$ and not the entire path. Denote by S the set of $(t_i, P(t_i))$ such that $\|x(t'_m) - P(t_i)\| \leq r$. The follow equation occurs:

$$f_m([\delta/\tau]) = \sum_{(t_i, P(t_i)) \in S \wedge (t'_m - t_i) \tau = [\delta/\tau]} e^{-\left(\frac{\|x(t'_m) - P(t_i)\|}{\sigma r}\right)^2} + f_{m-1}(\delta)(1 - \xi)^{t'_m - t'_{m-1}}$$

[0132] This is described in Algorithm 2.2 below. The term f is used as a sparse vector that
 15 receives also negative integer indices. The set S is the set of all neighbors of $x(t_i)$ on the path and can be computed quickly using the PPLEB algorithm. It is easy to verify that if the number of neighbors of $x(t_i)$ is bounded by some constant n_{near} then the number of non-zeros in the vector f is bounded by n_{near}/ξ which is only a constant factor larger. The final stage of the algorithm is to output a specific value of δ if $f([\delta/\tau])$ is above some threshold value.

20

Algorithm 2 Efficient likelihood update.

```

1:  $\Gamma \leftarrow 0$ 
2: while  $(t_j', x(t_j')) \in \text{INPUT}$  do
3:    $\Gamma \leftarrow (1 - \zeta_j)^{t_j' - t_j} \Gamma$ 
4:    $S \leftarrow \{(t_i, P(t_i)) \mid \|x(t_j') - P(t_i)\| \leq r\}$ 
5:   for  $(t_i, P(t_i)) \in S$  do
6:      $\delta \leftarrow t_j' - t_i$ 
7:
           
$$f(|\delta/\tau|) \leftarrow f(|\delta/\tau|) + e^{-\left(\frac{\|x(t_j') - P(t_i)\|}{\sigma r}\right)^2}$$

8:   end for
9:   Set all  $f$  values below threshold  $\epsilon$  to zero.
10: end while

```

[0133] FIG. 8 gives three consecutive point locations and the path points around them. Note that neither the bottom point nor middle one alone would have been sufficient to identify the correct part of the path. Together, however, they are. Adding the top point increases the certainty that the particle is indeed of the final (left) curve of the path.

[0134] In FIG. 9, given a set of n (grey) points, the algorithm is given a query point (black) and returns the set of points that lie within distance r from it (the points inside the circle). In the traditional setting, the algorithm must return all such points. In the probabilistic setting each such point should be returned only with some constant probability.

[0135] FIG. 10 illustrates the values of $u(x_1)$, $u(x_2)$, and $u(x)$. Intuitively, the function u gives different values to x_1 and x_2 if the dashed line passes between them and the same value otherwise. Passing the dashed line in a random direction ensures that the probability of this happening is directly proportional to angle between x_1 and x_2 .

[0136] FIG. 11 shows that by dividing the space into rings such that ring R_i is between radius $2r(1+\epsilon)^i$ and $2r(1+\epsilon)^{i+1}$, it can be made sure that any two vectors within a ring are the same length up to $(1+\epsilon)$ factors and that any search is performed in at most $1/\epsilon$ rings.

[0137] FIG. 12 shows a self-intersecting paths and a query point (in black). It illustrates that without the history of the particle positions it is impossible to know where it is on the path.

[0138] FIG. 13 gives three consecutive point locations and the path points around them. Note that neither $x(t_1)$ nor $x(t_2)$ alone would have been sufficient to identify the correct part of the path. Together however they are. Adding $x(t_3)$ increases the certainty that the particle is indeed of the final (left) curve of the path.

5 [0139] Substantial variations may be made in accordance with specific requirements. For example, customized hardware might also be used, and/or particular elements might be implemented in hardware, software (including portable software, such as applets, etc.), or both. Further, connection to other access or computing devices such as network input/output devices may be employed.

10 [0140] In the foregoing specification, aspects of the invention are described with reference to specific examples thereof, but those skilled in the art will recognize that the invention is not limited thereto. Various features and aspects of the above-described invention may be used individually or jointly. Further, examples can be utilized in any number of environments and applications beyond those described herein without departing from the broader spirit and scope
15 of the specification. The specification and drawings are, accordingly, to be regarded as illustrative rather than restrictive.

[0141] In the foregoing description, for the purposes of illustration, methods were described in a particular order. It should be appreciated that in alternate examples, the methods may be performed in a different order than that described. It should also be appreciated that the methods
20 described above may be performed by hardware components or may be embodied in sequences of machine-executable instructions, which may be used to cause a machine, such as a general-purpose or special-purpose processor or logic circuits programmed with the instructions to perform the methods. These machine-executable instructions may be stored on one or more machine readable mediums, such as CD-ROMs or other type of optical disks, floppy diskettes,
25 ROMs, RAMs, EPROMs, EEPROMs, magnetic or optical cards, flash memory, or other types of machine-readable mediums suitable for storing electronic instructions. Alternatively, the methods may be performed by a combination of hardware and software.

[0142] Where components are described as being configured to perform certain operations, such configuration can be accomplished, for example, by designing electronic circuits or other

hardware to perform the operation, by programming programmable electronic circuits (e.g., microprocessors, or other suitable electronic circuits) to perform the operation, or any combination thereof.

5 [0143] While illustrative examples of the application have been described in detail herein, it is to be understood that the inventive concepts may be otherwise variously embodied and employed, and that the appended claims are intended to be construed to include such variations, except as limited by the prior art.

0 [0144] Throughout this specification and the claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" or "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

5 [0145] The reference in this specification to any prior publication (or information derived from it), or to any matter which is known, is not, and should not be taken as, an acknowledgement or admission or any form of suggestion that that prior publication (or information derived from it) or known matter forms part of the common general knowledge in the field of endeavour to which this specification relates.

The claims defining the invention are as follows:

1. A computer-implemented method comprising:
 - receiving, by a server, a reference data set associated with a media segment, wherein the server is configured to identify an unidentified media segment by matching an unidentified data set with the reference data set, and wherein the unidentified data set is associated with the unidentified media segment;
 - storing the reference data set in a reference database;
 - receiving popularity data indicating a popularity of the media segment, wherein the popularity data includes values associated with at least one or more of viewing information of the media segment, rating information of the media segment, or information associated with a posted message on a remote source, and wherein the posted message is associated with the media segment;
 - determining a deletion time for the reference data set, wherein the deletion time is determined by calculating a representative value of two or more values of the popularity data, the representative value being one of an average, a mean, or a median of the two or more of the values of the popularity data; and
 - deleting the reference data set from the reference database, wherein the reference data set is deleted after the deletion time has expired.
2. The computer-implemented method of claim 1, further comprising:
 - adding the reference data set to a backup database when the reference data set is deleted from the reference database, wherein the server is configured to search the backup database after the reference database for a data set matching the unidentified data set.
3. The computer-implemented method of claim 1 or 2, wherein the reference data set includes pixel data or audio data for a frame of the media segment.
4. The computer-implemented method of any one of claims 1 to 3, wherein the information associated with the posted message on the remote source includes at least one or more of a number of posted messages on the remote source, a number of the posted messages that are positive, a number of the posted messages that are negative, a number of positive indications to the posted message, a number of reposts of the posted message, or a number of views of the posted message, and wherein the posted messages are associated with the media segment.

5. The computer-implemented method of any one of claims 1 to 4, wherein the viewing information includes a number of times the media segment has been identified by the server.

6. The computer-implemented method of any one of claims 1 to 5, wherein determining a deletion time for the reference data set comprises using a respective weight associated with one or more of the values of the popularity data.

7. A system comprising:
one or more processors; and
a non-transitory computer-readable medium containing instructions that, when executed by the one or more processors, cause the one or more processors to perform operations including:

receive, by a server, a reference data set associated with a media segment, wherein the server is configured to identify an unidentified media segment by matching an unidentified data set with the reference data set, and wherein the unidentified data set is associated with the unidentified media segment;

store the reference data set in a reference database;

receive popularity data indicating a popularity of the media segment, wherein the popularity data includes values associated with at least one or more of viewing information of the media segment, rating information of the media segment, or information associated with a posted message on a remote source, and wherein the posted message is associated with the media segment;

determine a deletion time for the reference data set, wherein the deletion time is determined by calculating a representative value of two or more values of the popularity data, the representative value being one of an average, a mean, or a median of the two or more of the values of the popularity data; and

delete the reference data set from the reference database, wherein the reference data set is deleted after the deletion time has expired.

8. The system of claim 7, further comprising instructions that, when executed by the one or more processors, cause the one or more processors to perform operations including:

add the reference data set to a backup database when the reference data set is deleted from the reference database, wherein the server is configured to search the backup database after the reference database for a data set matching the unidentified data set.

9. The system of claim 7 or 8, wherein the reference data set includes pixel data or audio data for a frame of the media segment.

10. The system of any one of claims 7 to 9, wherein the information associated with the posted message on the remote source includes at least one or more of a number of posted messages on the remote source, a number of the posted messages that are positive, a number of the posted messages that are negative, a number of positive indications to the posted message, a number of reposts of the posted message, or a number of views of the posted message, and wherein the posted messages are associated with the media segment.

11. The system of any one of claims 7 to 10, wherein the viewing information includes a number of times the media segment has been identified by the server.

12. The system of any one of claims 7 to 11, wherein determining a deletion time for the reference data set comprises using a respective weight associated with one or more of the values of the popularity data.

13. A computer-program product tangibly embodied in a non-transitory machine-readable storage medium, including instructions that, when executed by the one or more processors, cause the one or more processors to:

receive, by a server, a reference data set associated with a media segment, wherein the server is configured to identify an unidentified media segment by matching an unidentified data set with the reference data set, and wherein the unidentified data set is associated with the unidentified media segment;

store the reference data set in a reference database;

receive popularity data indicating a popularity of the media segment, wherein the popularity data includes values associated with at least one or more of viewing information of the media segment, rating information of the media segment, or information associated with a posted message on a remote source, and wherein the posted message is associated with the media segment;

determine a deletion time for the reference data set, wherein the deletion time is determined by calculating a representative value of two or more values of the popularity data,

the representative value being one of an average, a mean, or a median of the two or more of the values of the popularity data; and

delete the reference data set from the reference database, wherein the reference data set is deleted after the deletion time has expired.

14. The computer-program product of claim 13, further including instructions that, when executed by the one or more processors, cause the one or more processors to:

add the reference data set to a backup database when the reference data set is deleted from the reference database, wherein the server is configured to search the backup database after the reference database for a data set matching the unidentified data set.

15. The computer-program product of claim 13 or 14, wherein the reference data set includes pixel data or audio data for a frame of the media segment.

16. The computer-program product of any one of claims 13 to 15, wherein the information associated with the posted message on the remote source includes at least one or more of a number of posted messages on the remote source, a number of the posted messages that are positive, a number of the posted messages that are negative, a number of positive indications to the posted message, a number of reposts of the posted message, or a number of views of the posted message, and wherein the posted messages are associated with the media segment.

17. The computer-program product of any one of claims 13 to 16, wherein determining a deletion time for the reference data set comprises using a respective weight associated with one or more of the values of the popularity data.

18. The computer-program product of any one of claims 13 to 17, wherein the viewing information includes a number of times the media segment has been identified by the server.

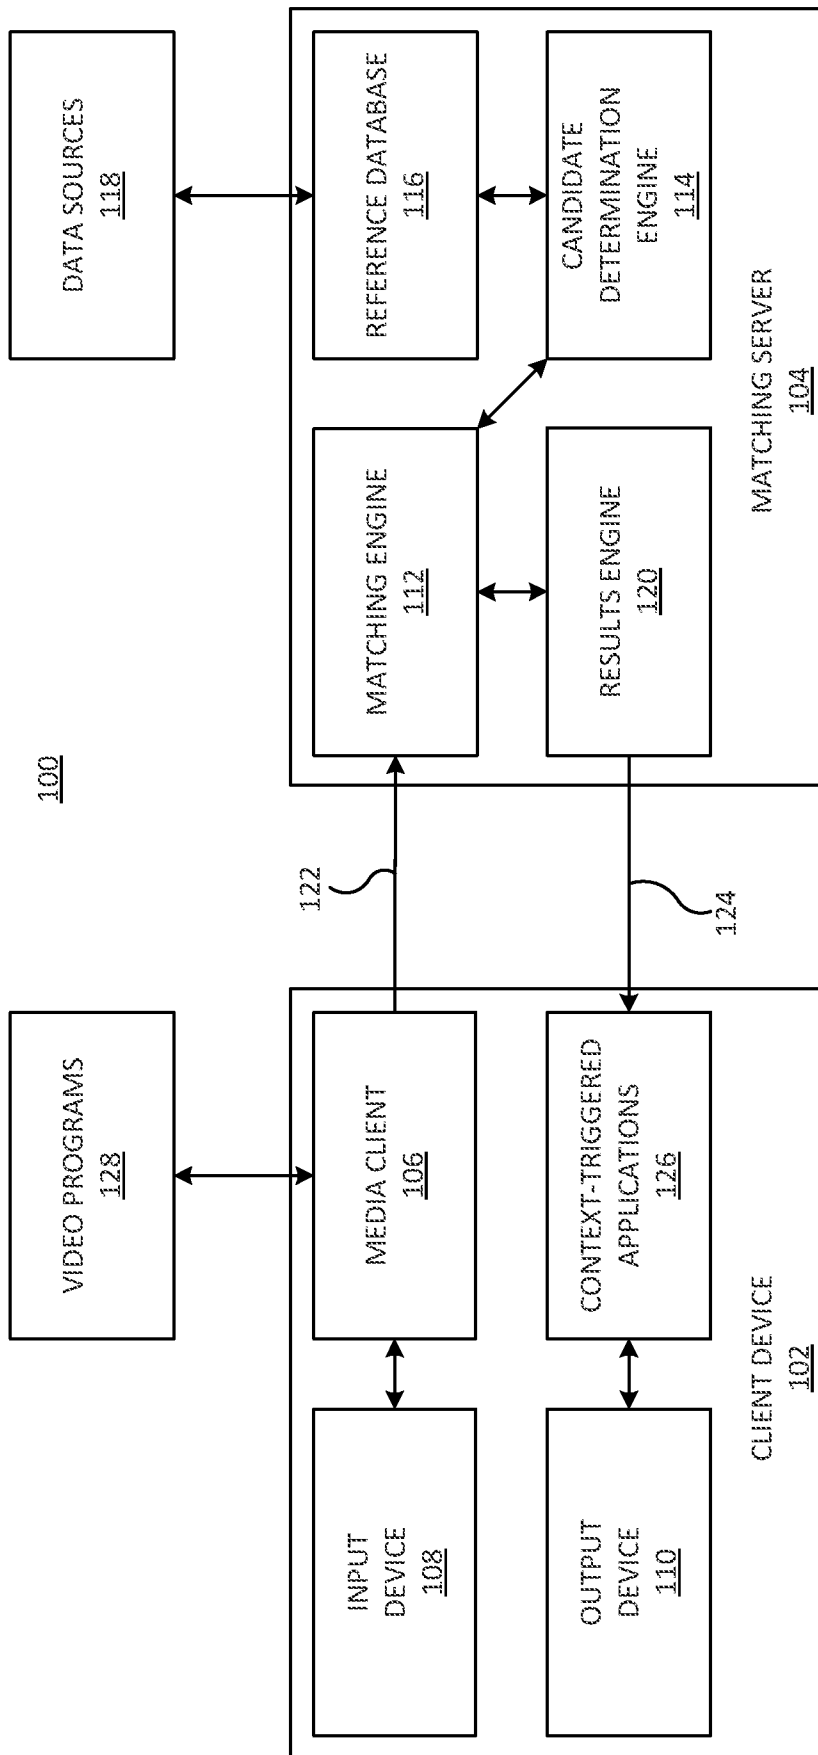


FIG. 1

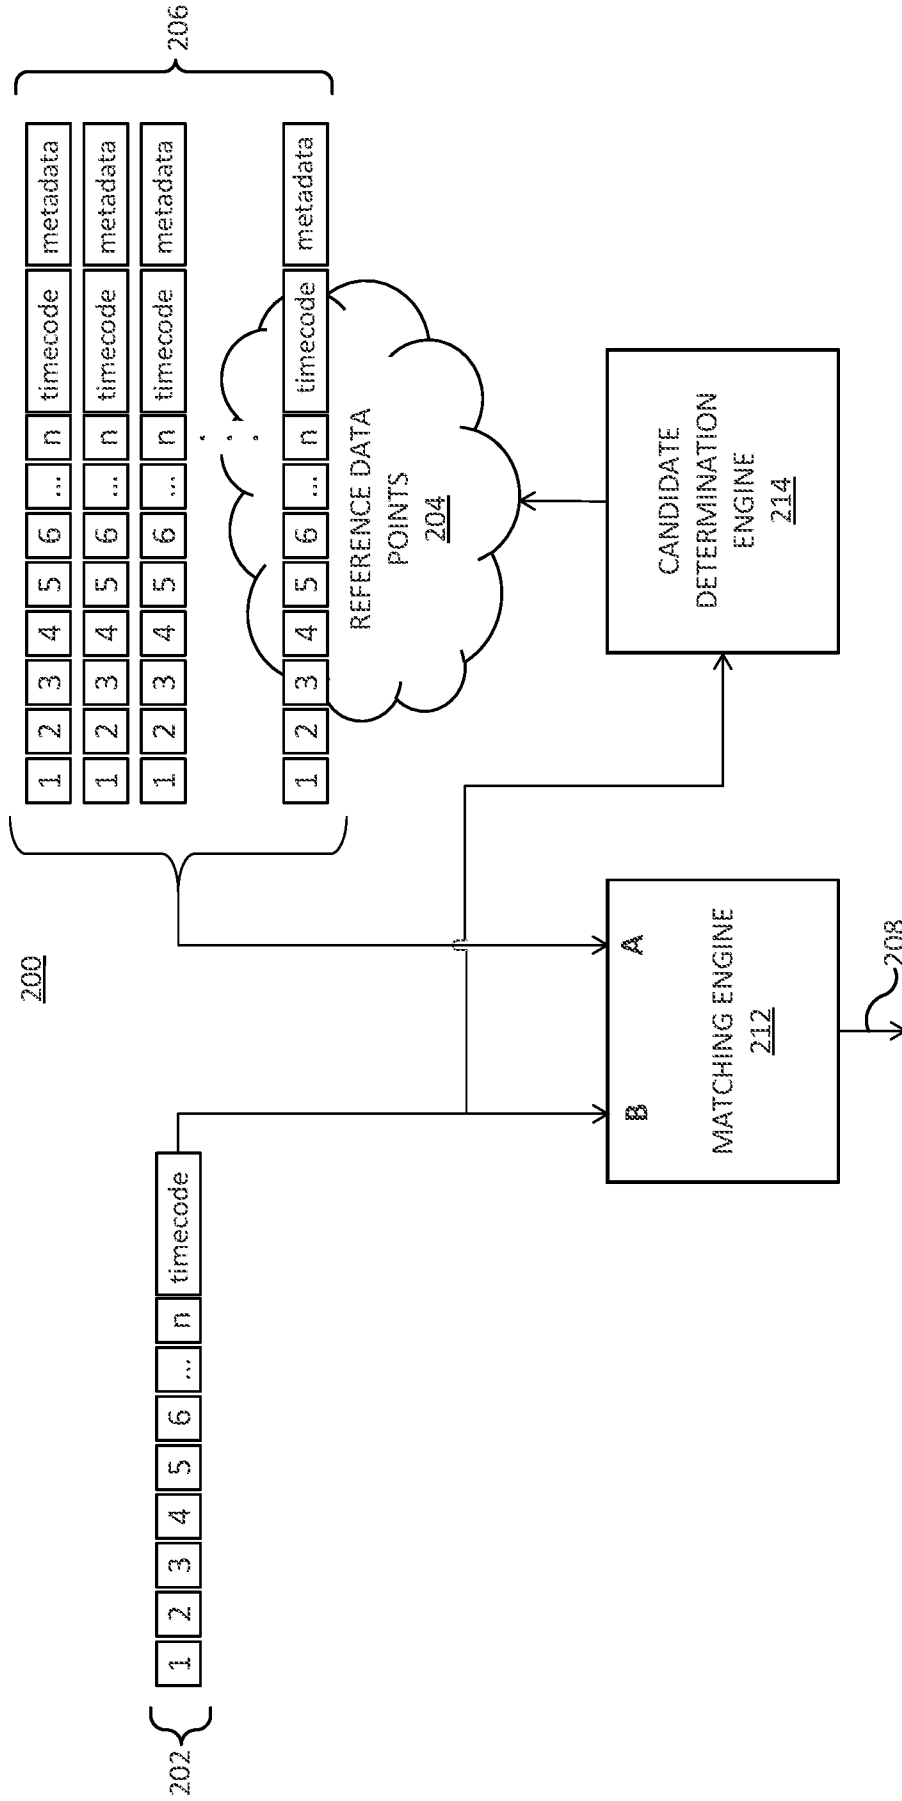


FIG. 2

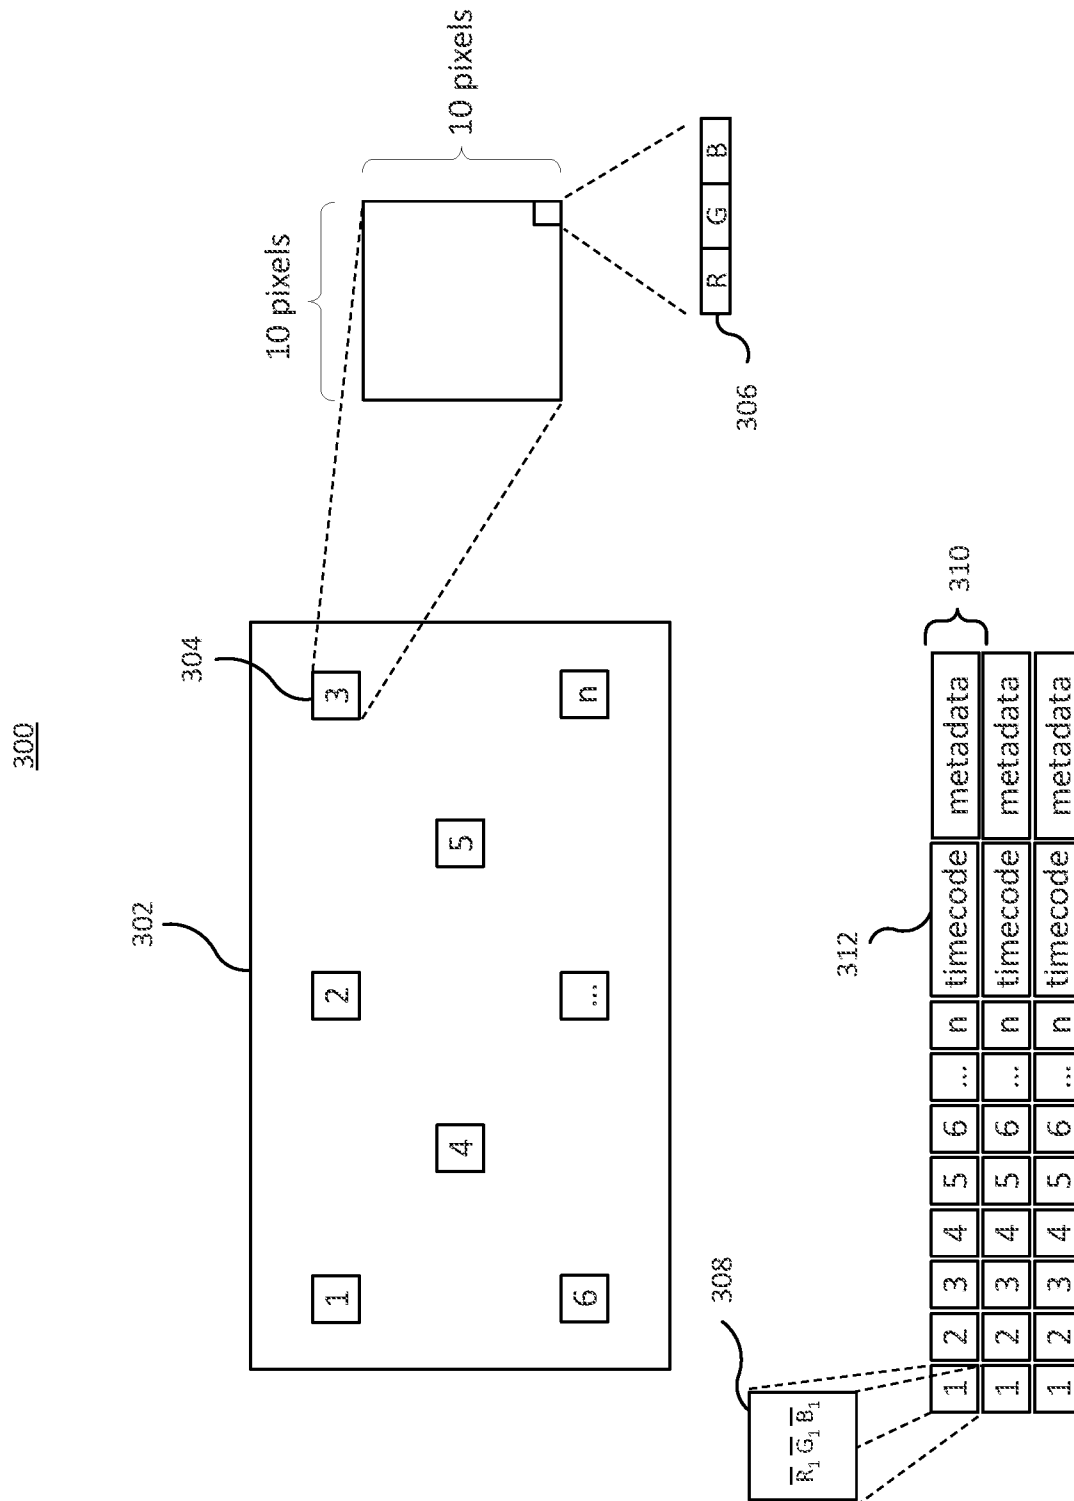


FIG. 3

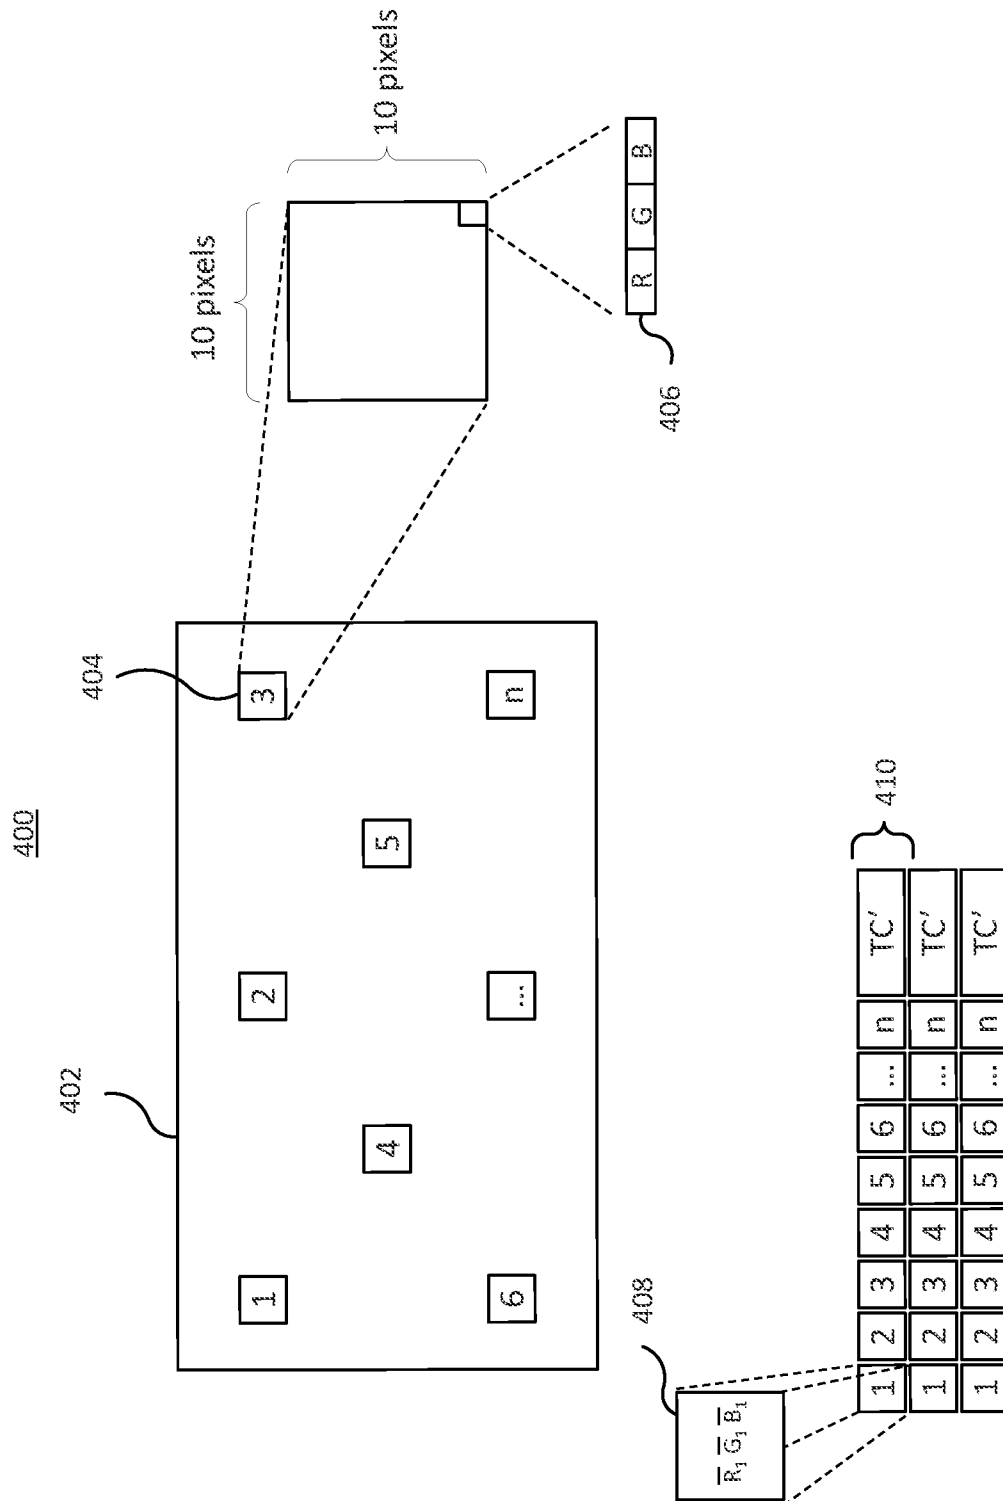
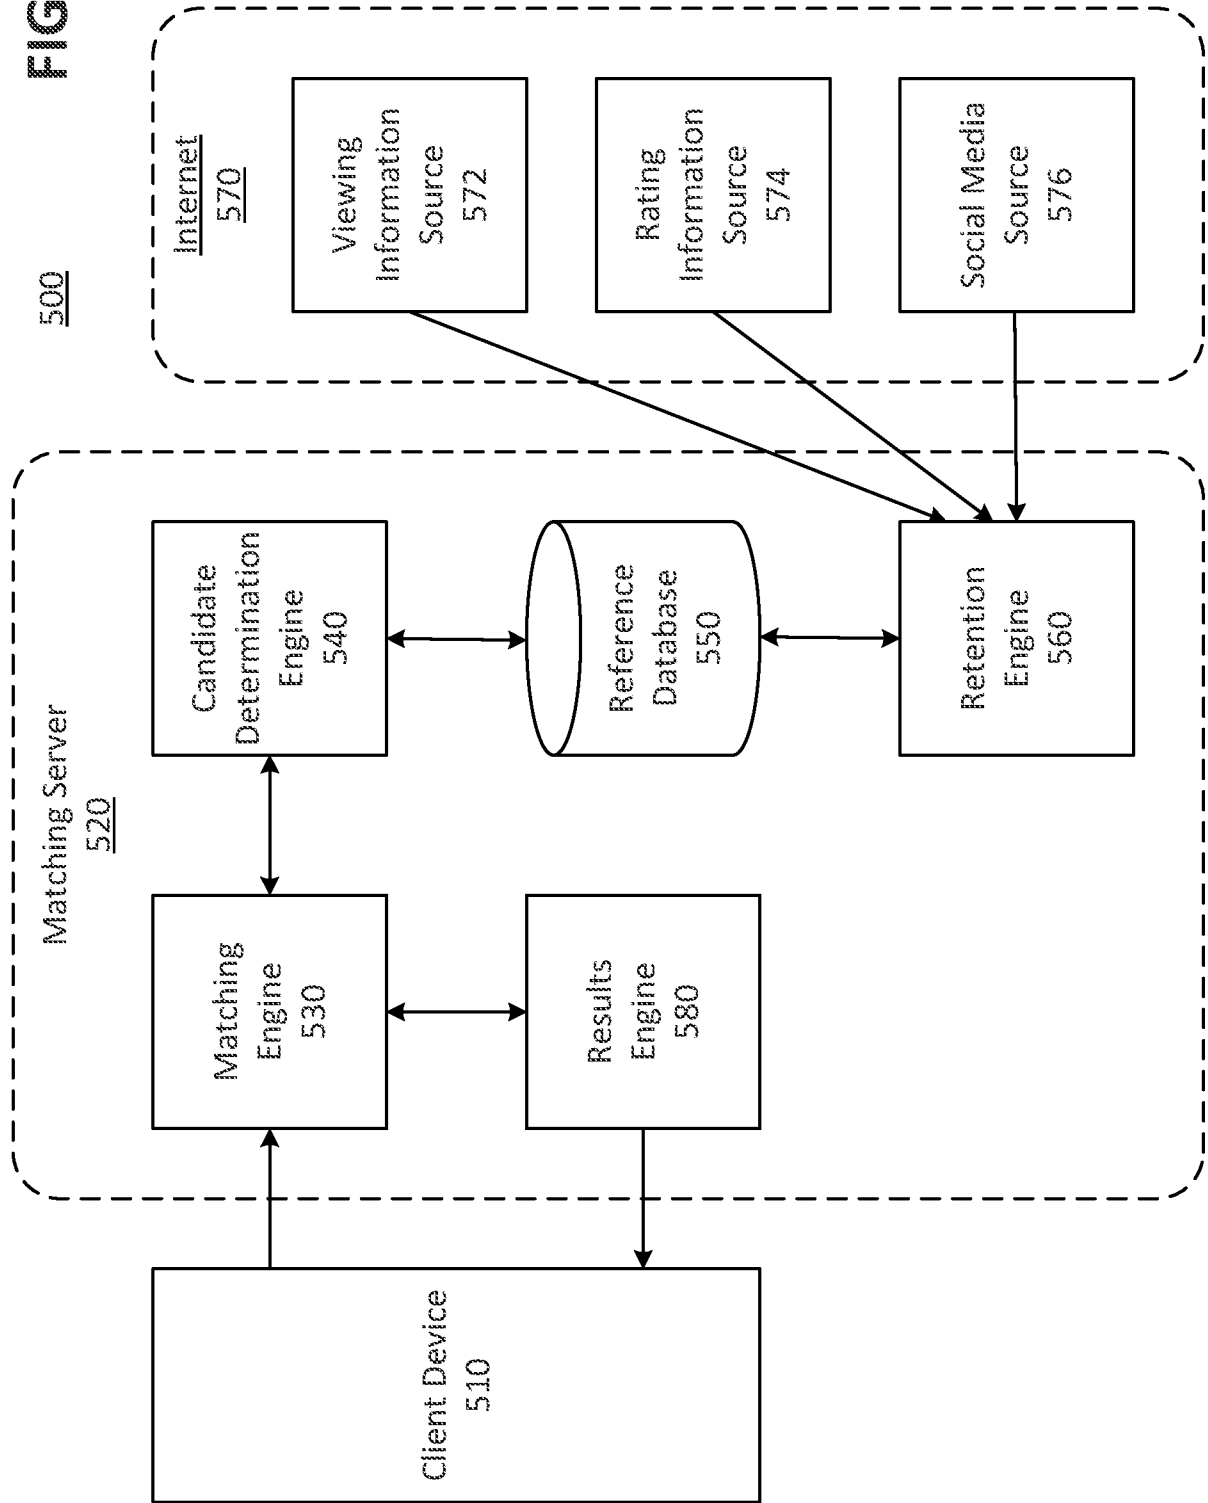


FIG. 4

FIG. 5



600

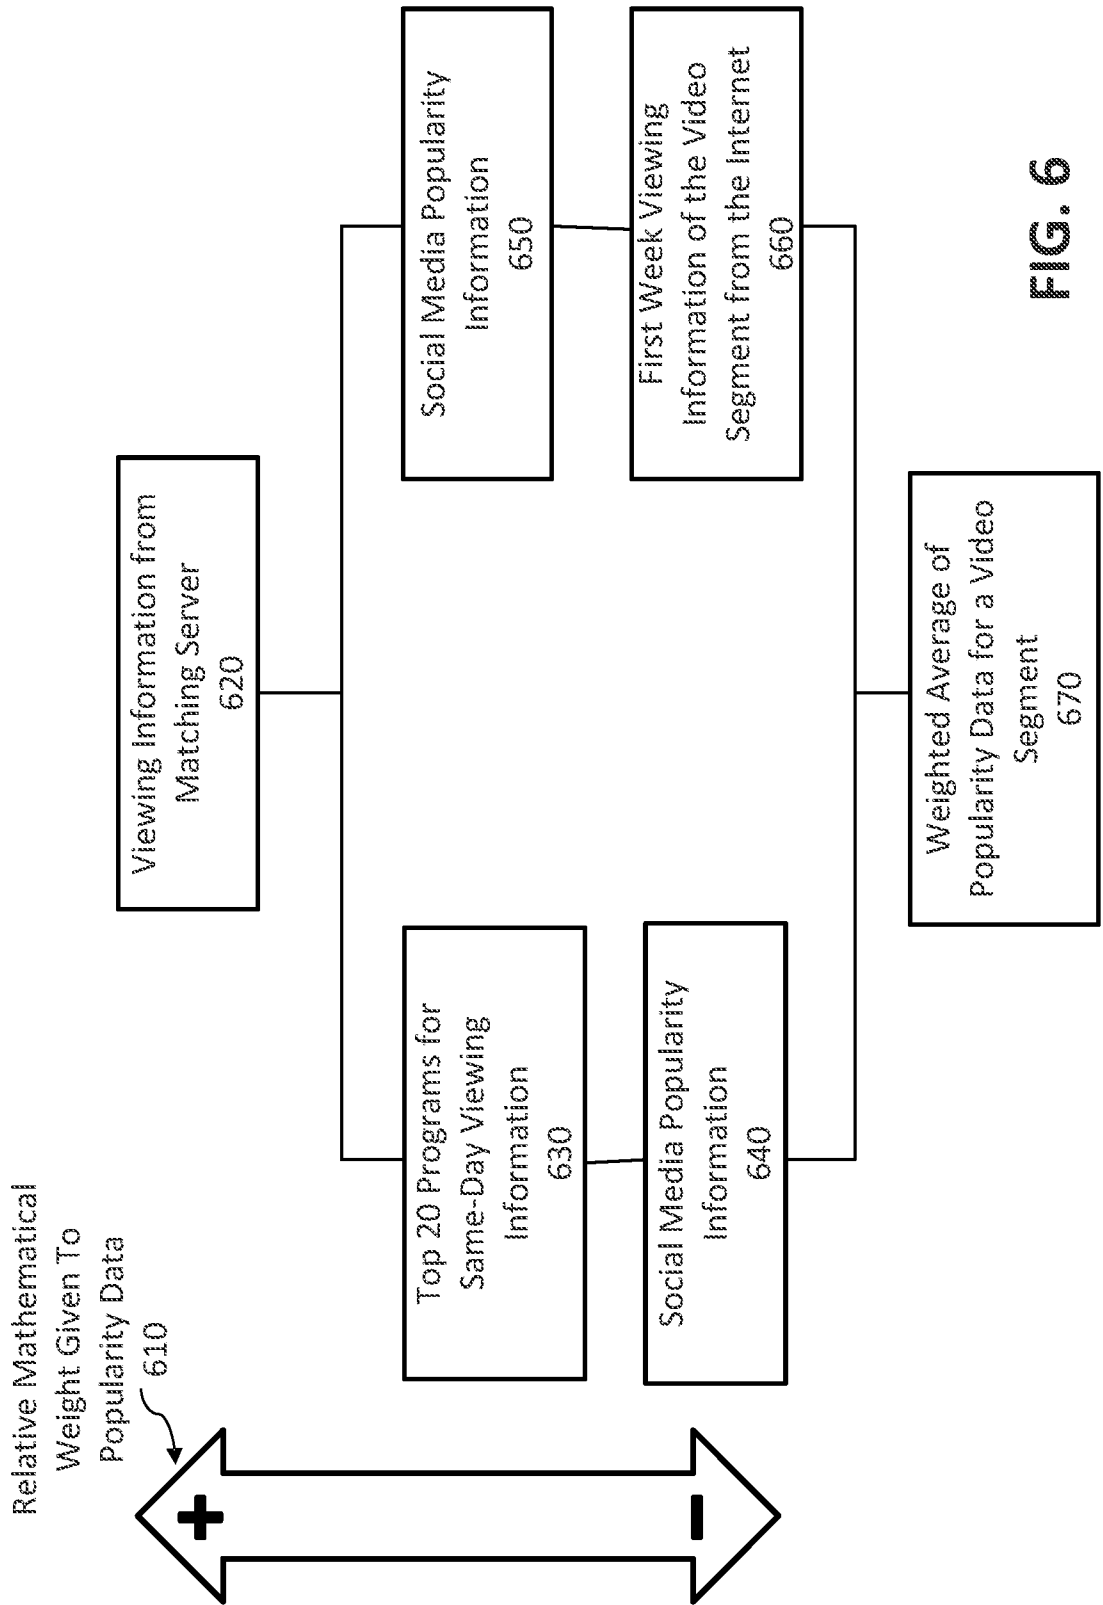
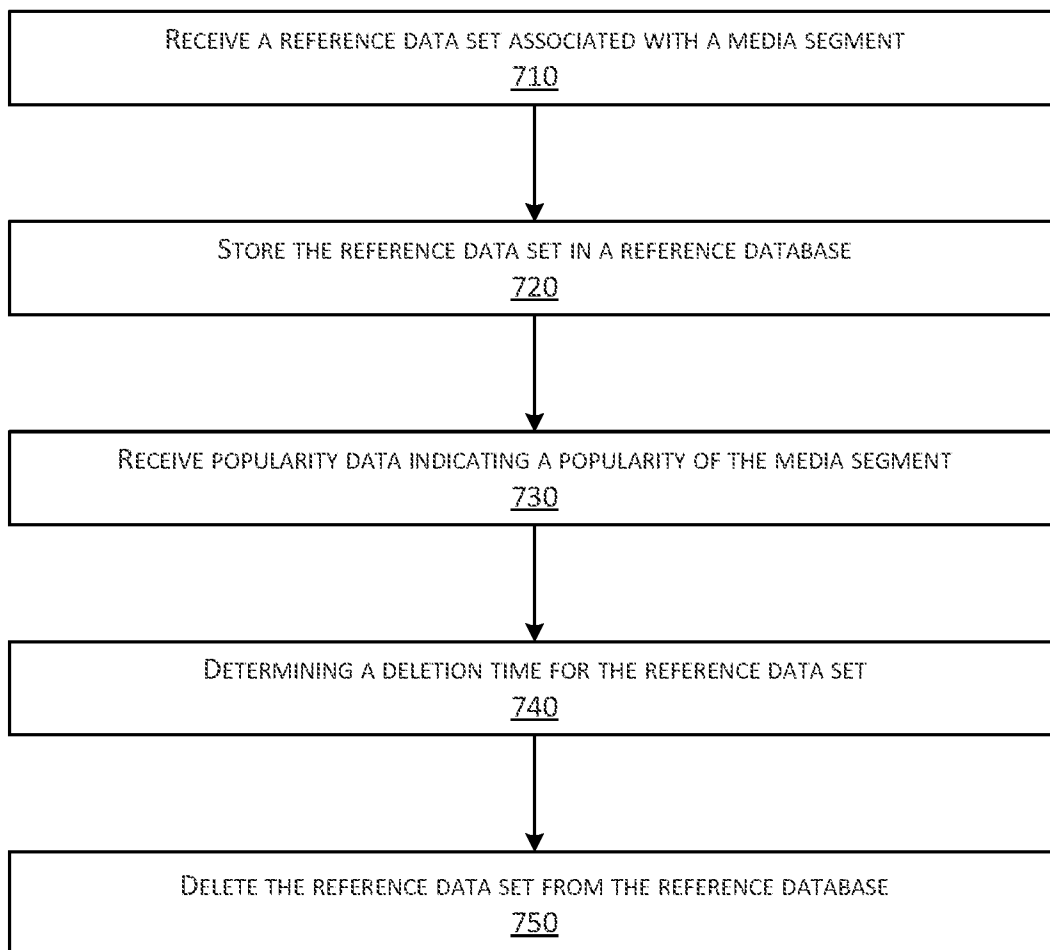


FIG. 6

7/13

700**FIG. 7**

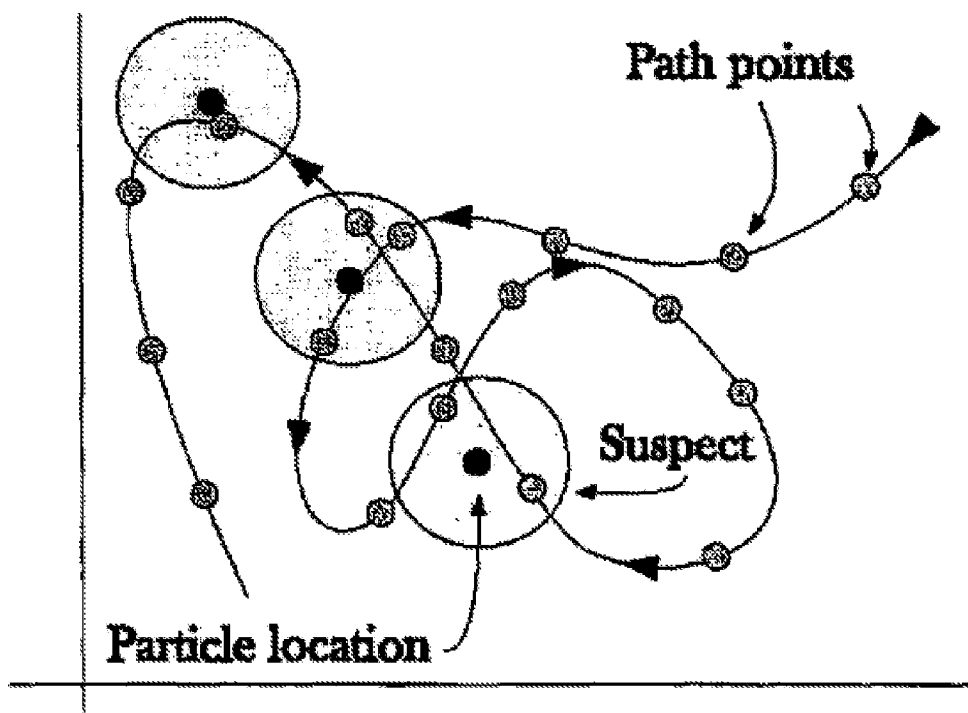


FIG. 8

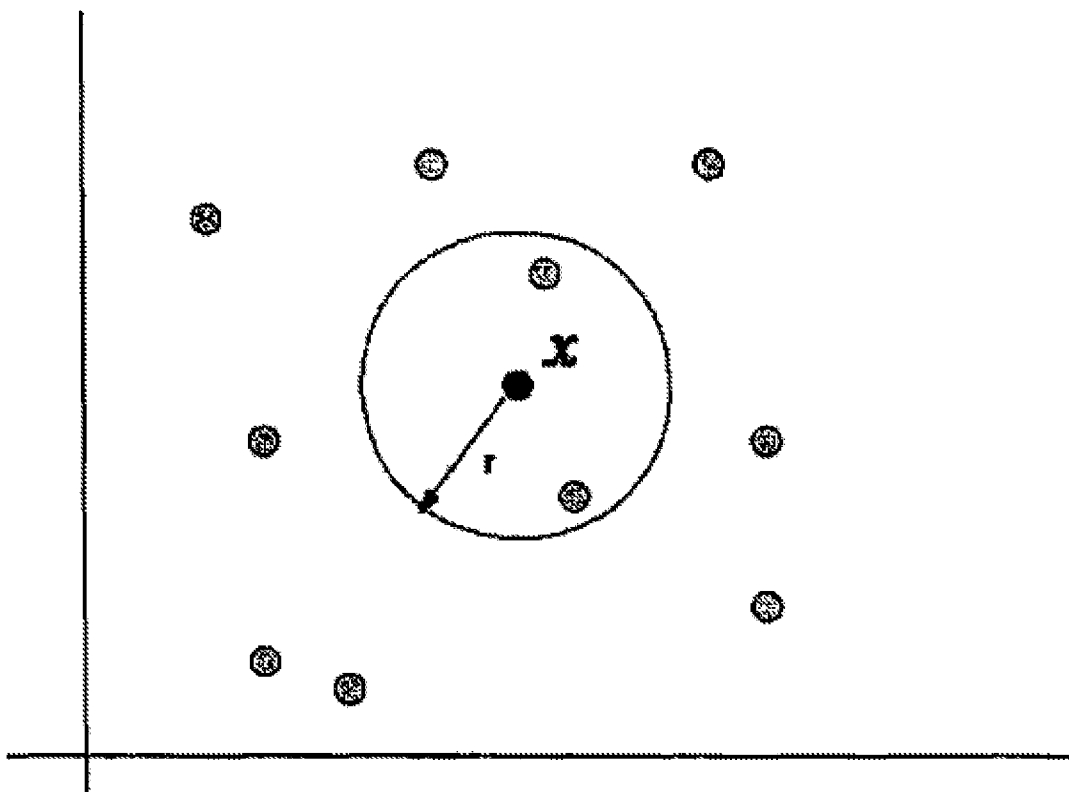


FIG. 9

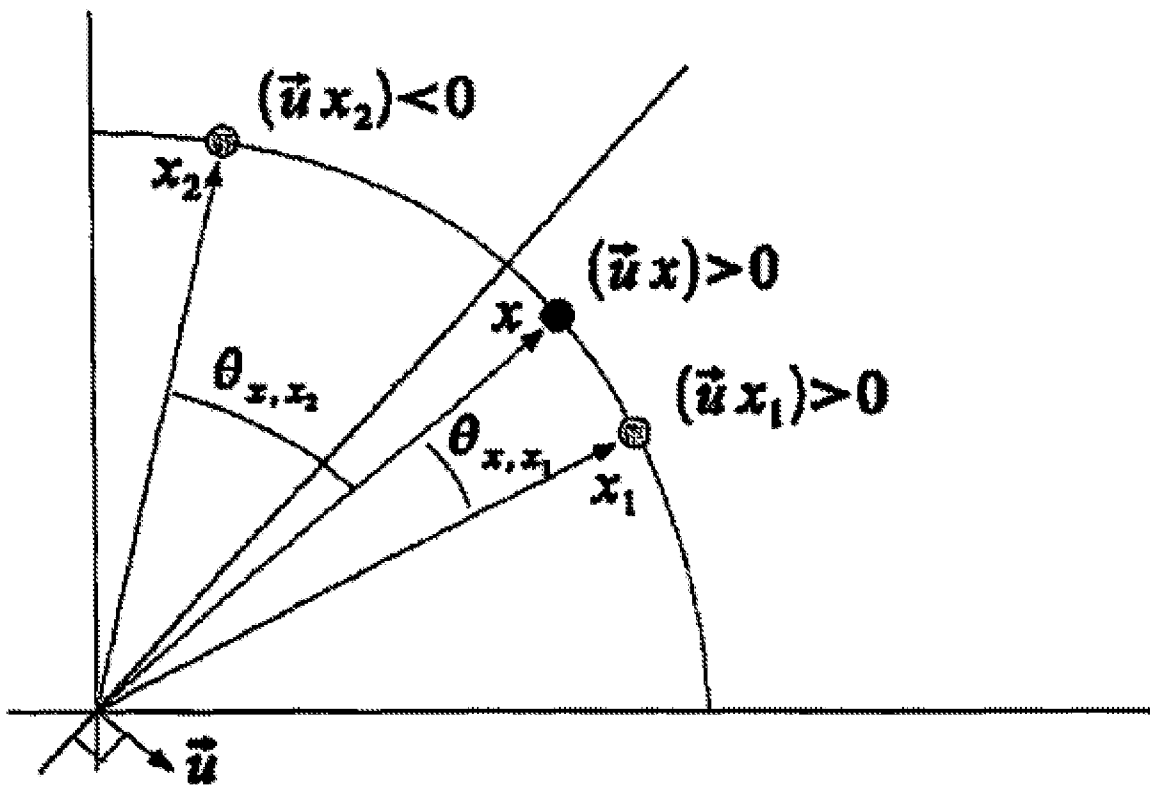


FIG. 10

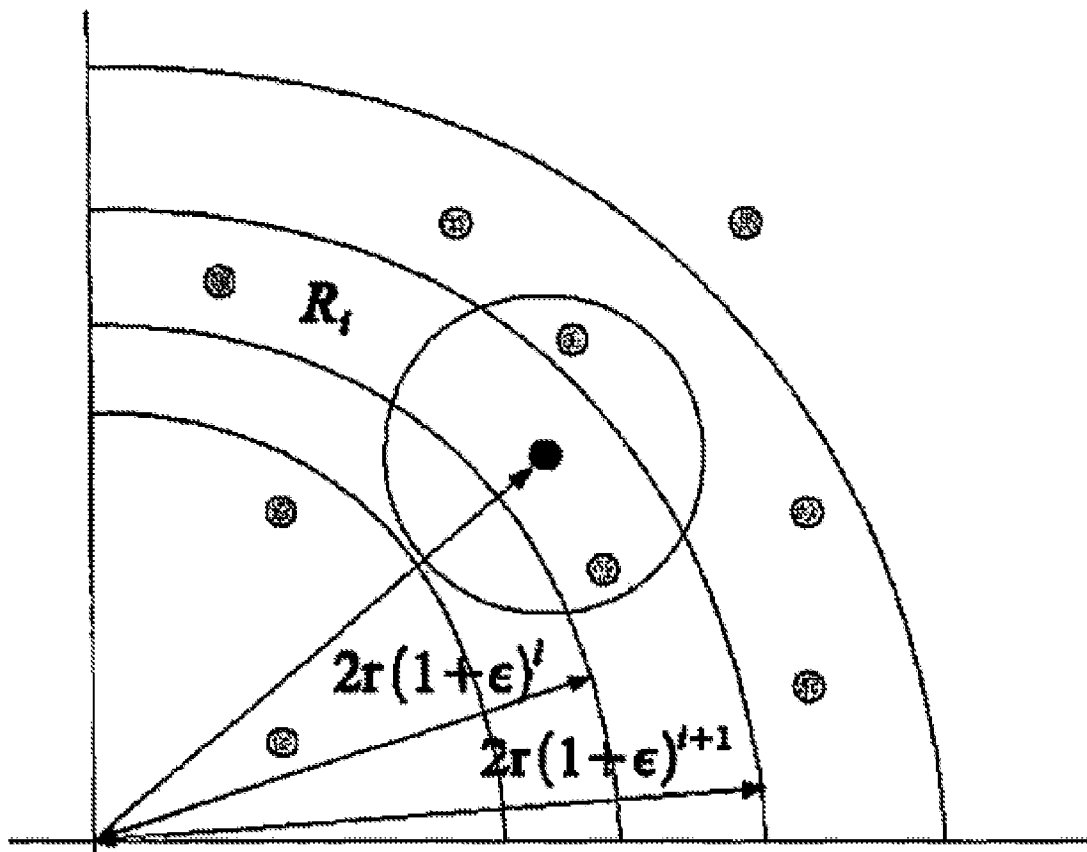


FIG. 11

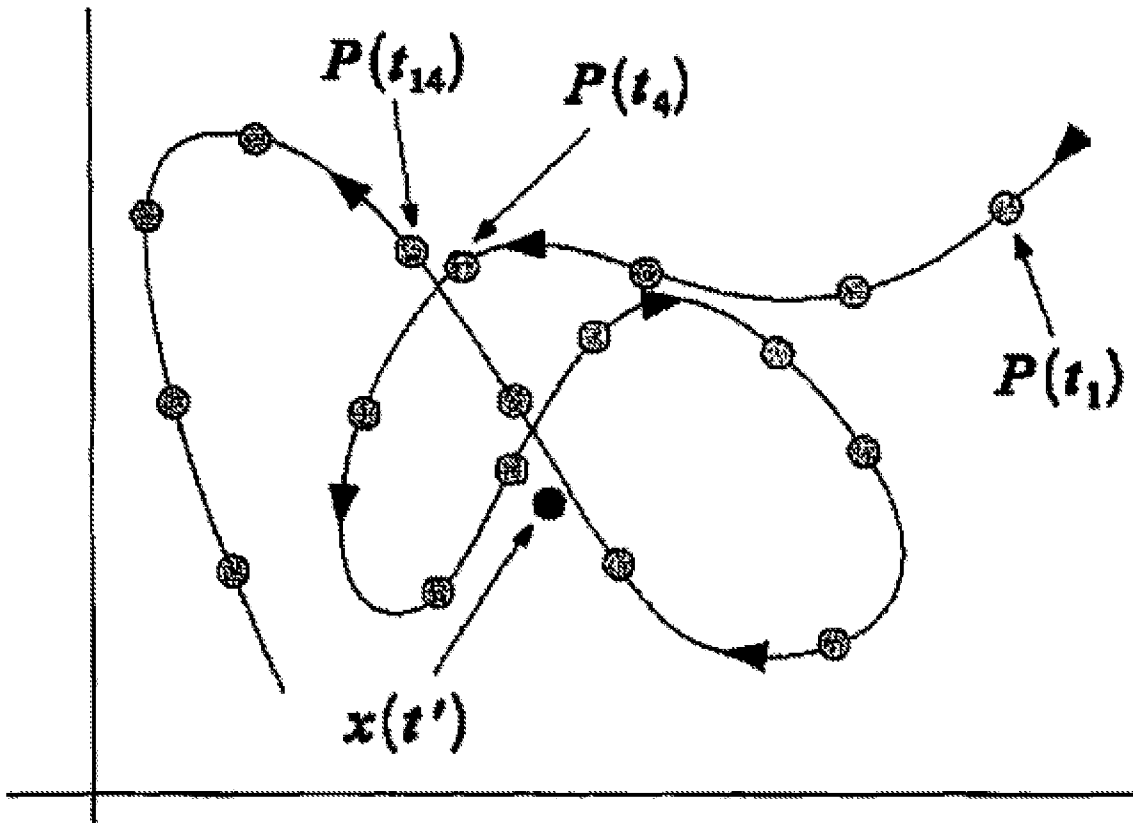


FIG. 12

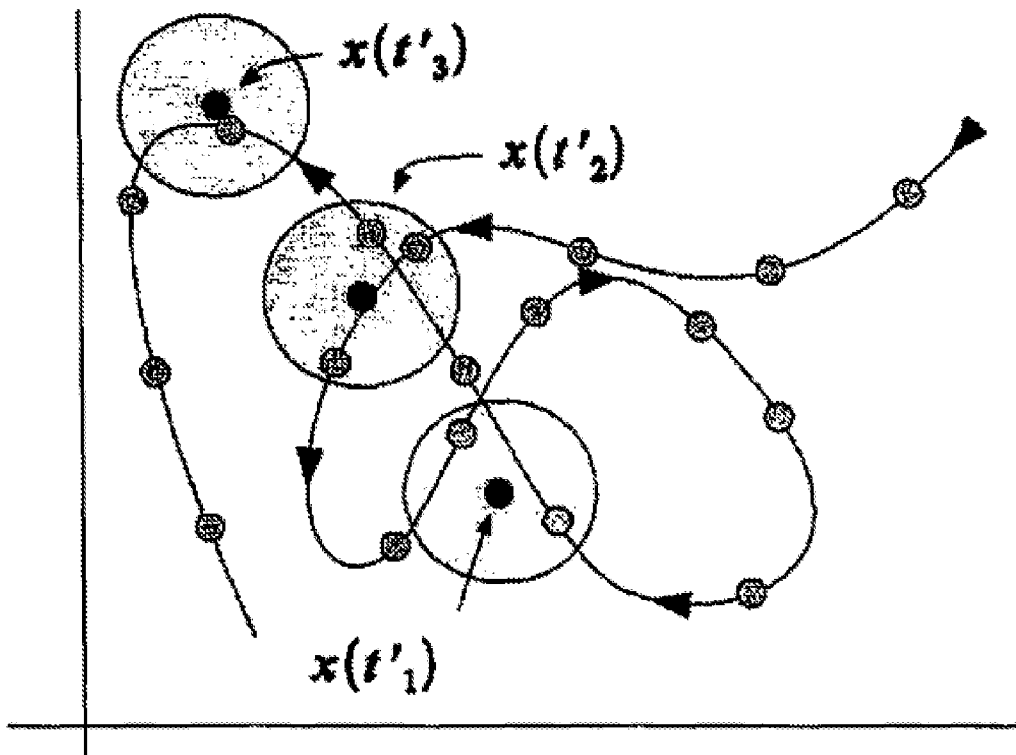


FIG. 13