

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2017249085 B2**

(54) Title
Generation, ranking, and delivery of actions for entities in a video delivery system

(51) International Patent Classification(s)
H04N 21/442 (2011.01) **H04N 21/45** (2011.01)

(21) Application No: **2017249085** (22) Date of Filing: **2017.04.05**

(87) WIPO No: **WO17/180396**

(30) Priority Data

(31) Number	(32) Date	(33) Country
15/399,702	2017.01.05	US
62/323,235	2016.04.15	US

(43) Publication Date: **2017.10.19**

(44) Accepted Journal Date: **2020.11.26**

(71) Applicant(s)
Hulu, LLC

(72) Inventor(s)
Kaya, Lutfi Ilke;Tang, Bangsheng;Yang, Tong;Kehler, Christopher Russell;Zhang, Chi

(74) Agent / Attorney
Shelston IP Pty Ltd., L 9 60 Margaret St, Sydney, NSW, 2000, AU

(56) Related Art
US 2010/0211636 A1
US 2012/0204201 A1

(51) International Patent Classification:
H04N 21/442 (2011.01) *H04N 21/45* (2011.01)(21) International Application Number:
PCT/US2017/026195(22) International Filing Date:
5 April 2017 (05.04.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/323,235 15 April 2016 (15.04.2016) US
15/399,702 5 January 2017 (05.01.2017) US

(71) Applicant: HULU, LLC [US/US]; 2500 Broadway, Suite 200, Santa Monica, California 90404 (US).

(72) Inventors: KAYA, Lutfi Ilke; 2500 Broadway, Suite 200, Santa Monica, California 90404 (US). TANG, Bangsheng; 2500 Broadway, Suite 200, Santa Monica, California 90404 (US). YANG, Tong; 2500 Broadway, Suite 200, Santa Monica, California 90404 (US). KEHLER, Christopher Russell; 2500 Broadway, Suite 200, Santa Monica, California 90404 (US). ZHANG, Chi; 2500 Broadway, Suite 200, Santa Monica, California 90404 (US).

(74) Agent: YOUNG, Brian N.; Fountainhead Law Group P.C., 900 Lafayette St., Suite 301, Santa Clara, California 95050 (US).

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

[Continued on next page]

(54) Title: GENERATION, RANKING, AND DELIVERY OF ACTIONS FOR ENTITIES IN A VIDEO DELIVERY SYSTEM




FIG. 6A

(57) Abstract: In one embodiment, a method sends videos to a user that is using a video delivery service. The method receives user behavior that includes actions taken by the user on the video delivery service. The method inputs the user behavior into a first predictor to generate a set of actions for a set of entities. Also, the method inputs the set of actions for the set of entities, a real-time context, and the user behavior into a second predictor to generate probabilities for the set of actions for the set of entities. A probability for an action indicates the probability the user would select that action for an entity when compared against other actions in the set of actions for the set of entities. An action feed is selected based on the ranking and dynamically output to a client while the user is using the video delivery service.



Published:

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

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

Generation, Ranking, and Delivery of Actions for Entities in a Video Delivery System

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] The present application also claims priority to U.S. Nonprovisional Application No. 15/399,702, entitled “Generation, Ranking, and Delivery of Actions for Entities in a Video Delivery System”, filed January 5, 2017, which claims priority to U.S. Provisional App. No. 62/323,235, entitled “Contextual Generation of Actions for a Watch List of Entities in a Video Delivery System,” filed April 15, 2016. The entire contents of these applications are incorporated herein by reference in their entireties for all purposes.

BACKGROUND

[0002] Television (TV) broadcast has been a medium where users were responsible for tracking content and scheduling their lives around the schedule of shows on the TV broadcast. As a result of this reactive nature of the TV broadcast, users miss content that they would not miss otherwise if they were given a chance to watch it. Some reasons for missing content can be the user not being available to watch the content at a certain time (and having to manually program a digital video recorder (DVR) to record the content), or not knowing that the content is being broadcasted or is available for only a short period of time. Also, in the case that the content is indefinitely available – simply not knowing that the content exists. For example, a user may have a favorite actor, and that actor may appear as a guest actor in a single episode of a TV show. The user may never have watched a TV show, and may not be interested in watching the whole TV show at the moment, but it is possible the user would watch that one episode and see how the actor performed if the user knew of the episode. With how current TV broadcast works, the user would have to find out independently that the actor appeared in the episode, search for the episode, and then record or watch the episode when the episode is broadcasted at a certain time. The user always has to know the schedule, manage the schedule, and track events.

[0003] Any discussion of the prior art throughout the specification should in no way be considered as an admission that such prior art is widely known or forms part of common general knowledge in the field.

[0004] It is an object of the present invention to overcome or ameliorate at least one of the disadvantages of the prior art, or to provide a useful alternative.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] FIG. 1 depicts a simplified system for generating and providing an action feed according to one embodiment.

[0006] FIG. 2 depicts a more detailed example of an action feed generator according to one embodiment.

[0007] FIG. 3 shows a graph of entities and the relationships between the entities according to one embodiment.

[0008] FIG. 4 depicts an example of an entity-to-content matching process according to one embodiment.

[0009] FIG. 5A depicts an example of the ranking process used by an action ranker according to one embodiment.

[0010] FIG. 5B depicts a table of entities and actions that have been ranked according to one embodiment.

[0011] FIG. 6A depicts a more detailed example of the action feed generator for displaying an action feed according to one embodiment.

[0012] FIG. 6B depicts an example of a table that is used to store the selected actions according to one embodiment.

[0013] FIG. 7 depicts an example of the action feed being displayed on a user interface of a client according to one embodiment.

[0014] FIG. 8 depicts a simplified flowchart of a method for generating an action feed according to one embodiment.

[0015] FIG. 9 depicts a simplified flowchart of a method for displaying the action feed according to one embodiment.

[0016] FIG. 10 depicts a simplified flowchart of a method for responding to a selection of an action in the action feed according to one embodiment.

[0017] FIG. 11 depicts a video streaming system in communication with multiple client devices via one or more communication networks according to one embodiment.

[0018] Fig. 12 depicts a diagrammatic view of an apparatus for viewing video content and advertisements.

DETAILED DESCRIPTION

[0019] Described herein are techniques for a system to generate actions for entities. In the following description, for purposes of explanation, numerous examples and specific details are set forth in order to provide a thorough understanding of particular embodiments. Particular embodiments as defined by the claims may include some or all of the features in these examples alone or in combination with other features described below, and may further include modifications and equivalents of the features and concepts described herein.

[0020] Particular embodiments generate and rank actions for entities instead of just suggesting content on a video delivery service. An action feed is generated that includes actions for entities that have been ranked against other actions for entities. Then, at least a portion of the actions can be displayed to a user via a user interface. For example, entities can be television (TV) shows, movies, people, genres, topics, sports teams, events and so on. Actions may include watch-specific actions and non-watch-specific actions, such as watch actions, follow actions, try actions, etc. Some examples of watch actions are “Start Watching Show #1”, “Resume Watching Show #2”, “Watch Trending Clip from Show #3”, and “Watch a Sports Game Live”. Some examples of non-watch actions include “Remind Me to Watch Event #1”, “Follow Actor #1”, and “Try No Commercial Add-on Package”. Particular embodiments may use a real-time context for a user and generate a ranking of actions for entities that are suggested to a user to perform at that current time. Then, at least some of the actions are provided to the user.

[0021] According to a first aspect of the present invention there is provided a method comprising: receiving, by a computing device, user behavior for a user, wherein the user behavior includes actions taken by the user on a video delivery service; receiving, by the computing device, a real-time context based on the user using the video delivery service;

generating, by the computing device, a set of action/entity pairs from a set of actions for a set of entities found on the video delivery service based on the user behavior, wherein each action/entity pair includes an action to be performed on an entity in which an action can be taken on the video delivery system; generating, by the computing device, probabilities for the set of action/entity pairs based on the real-time context, wherein a probability for an action/entity pair indicates the probability the user would select that action/entity pair when compared against other action/entity pairs; ranking, by the computing device, action/entity pairs in the set of action/entity pairs based on respective probabilities for the action/entity pairs; selecting, by the computing device, a list based on the ranking of the set of action/entity pairs, the list including at least a portion of the set of action/entity pairs; and outputting, by the computing device, the list to a client, wherein an action on an entity in the list is performed when selected by the user.

[0022] According to a second aspect of the present invention there is provided a non-transitory computer-readable storage medium containing instructions, that when executed, control a computer system to be configured for: receiving user behavior for a user, wherein the user behavior includes actions taken by the user on a video delivery service; receiving a real-time context based on the user using the video delivery service; generating a set of action/entity pairs from a set of actions for a set of entities found on the video delivery service based on the user behavior, wherein each action/entity pair includes an action to be performed on an entity found on the video delivery system; generating probabilities for the set of action/entity pairs based on real-time context, wherein a probability for an action/entity pair indicates the probability the user would select that action/entity pair when compared against other action/entity pairs; selecting a list based on the ranking of the set of action/entity pairs, the list including at least a portion of the set of action/entity pairs; and outputting the list to a client, wherein an action on an entity in the list is performed when selected by the user.

According to a third aspect of the present invention there is provided a method comprising: sending, by a computing device, videos to a user that is using a video delivery service on a client; receiving, by the computing device, user behavior for the user, wherein the user behavior includes actions taken by the user on the video delivery service; ranking first actions for manually added entities and second actions for engaged entities into a set of

third actions, the manually added entities being explicitly engaged by the user in the user behavior and the engaged entities being implicitly engaged by the user in the user behavior; and ranking fourth actions for discovered entities other than the engaged entities and manually added entities, the discovered entities being determined based a relevance between the engaged entities and the different entities; ranking fifth actions for a set of campaigns, the set of campaigns promoting campaign entities; ranking the ranked set of third actions and the ranked fourth actions into a set of sixth actions; selecting, by the computing device, a list based on the ranking of the set of sixth actions and the ranked fifth actions, the list including at least a portion of the set of actions for the set of entities; and dynamically outputting, by the computing device, the list to the client while the user is using the video delivery service, wherein an action on an entity in the list is performed when selected by the user.

[0023] Unless the context clearly requires otherwise, throughout the description and the claims, the words “comprise”, “comprising”, and the like are to be construed in an inclusive sense as opposed to an exclusive or exhaustive sense; that is to say, in the sense of “including, but not limited to”.

[0024] Overview

[0025] FIG. 1 depicts a simplified system 100 for generating and providing an action feed according to one embodiment. System 100 includes a server system 102 and multiple clients 104-1 to 104-N. Other components of a video delivery network may also be appreciated, such as content delivery networks, but are not shown. Also, server system 102 may include multiple servers and any number of clients 104 may be used.

[0026] A video delivery service may use a video delivery system 106 to provide content to users using client devices 104. Although a video delivery service is described, the video delivery service may provide content other than videos. For example, the video delivery service may include a library of media programs (e.g., videos, audio, or other forms of streamable content). Users may watch media programs that are available in the library through video delivery system 106 at any time. The library includes different kinds of media programs, such as movies, shows, shorts, clips, trailers, etc. The shorts, trailers, and clips may be shorter versions or previews of movies or shows. Movies are released once whereas shows include episodes that may be released (or initially aired) on a set schedule during seasons. For example, the multiple episodes for the shows may be released daily, weekly, or

on another schedule. Typically, the shows may be released seasonally, that is, a set number of episodes for a show may be released during a season, such as a four-month season during which an episode is released every week.

[0027] In the video delivery service, the content that a user can request is being refreshed regularly. For example, shows may include episodes that, when in season, are released onto the video delivery service weekly, daily, or in other intervals. This means the content is not stale. Although video delivery system 106 may provide shows that release episodes on a schedule, video delivery system 106 allows users to request the shows on-demand, such as once an episode for the show is released, the user may watch the episode at any time. This may be different from a user watching shows on a cable television system. In a cable television system, once the episode of the show is aired, the user cannot watch the show on demand unless it is offered on demand or the user recorded the show using the DVR. However, the cable television system may not receive an indication when a user watches the recorded shows on the DVR. But, for the video delivery system, the user needs to request the show from the service to watch the show. Also, a system that releases only single videos (e.g., movies) or all episodes for a show at once does not have the recurring timing information of when a user views released episodes week after week (or any other release schedule). Accordingly, the video delivery service faces a unique situation when releasing episodes of a show and also being able to record when a user watched the episodes. This allows the video delivery service to read contextual actions of the user with respect to entities (e.g., videos) on the video delivery service. Also, the videos released may be from a cable broadcast from content sources that then provide the videos to the video delivery service for release. Further, video delivery system 106 may provide linear television via the same service. The linear television may be similar to the cable television system described above. However, the linear television options can be blended with on demand options in an interface 112.

[0028] In one embodiment, the availability of both on demand and live options may complicate the recommendation process. That is, there may be different options for a user to watch content, such as there may be a live version and an on demand version of the same content. Additionally, a third option where a user records some content (like a digital video recorder) may also be available. Particular embodiments solve this problem by ranking

actions for entities instead of just determining what content is relevant. This allows recommendations for content that can be delivered differently to be ranked together.

[0029] An action feed generator 108 may generate an action feed 110 that includes actions for entities that each user may display on user interfaces 112-1 – 112-N. An entity may be something in which an action can be performed in the video delivery service. The entity may be something that exists as itself, as a subject or as an object, actually or potentially, concretely or abstractly, physically or virtually. For example, a TV show may be a video entity; an actor, director, or another person associated with the TV show may be a person entity; an adult animation may be a genre entity; a specific movie may be a movie entity; a specific user may be a user entity, and so on.

[0030] Action feed generator 108 uses a ranking of actions for entities against each other to determine a personalized order of actions for entities for action feed 110. Action feed generator 108 may then provide a personalized action feed 110 to each user on each user interface 112. Action feed 110 may include actions in a predicted order that action feed generator 108 thinks the user will want to perform at a given time.

[0031] To generate the actions, action feed generator 108 detects patterns in historical user behavior on the video delivery service, and ranks actions that fit into this behavior higher. As a result, action feed generator 108 can recommend actions to users as the actions are available, and not just full TV shows and movies that are related with the past shows and movies the user has watched.

[0032] Action feed generator 108 can display suggested actions to the user in several different categories or areas of user interface 112. When user interface 112 is discussed, it will be understood that this may be the interface displaying pages for an application, the application itself, and/or the physical interface on client 102. When a page is referred to, it will be recognized that this could be a section of the page. There can be pages in user interface 112 that contain categories of entities. When the user navigates to the page, a suggested action for each entity can be presented. The user can then take the action with a selection, such as a single selection, or the system can automatically take the action itself – starting playing without user engagement. In another example, the core surface point can be a dedicated area of user interface 112. When users turn on their client device 104 for the video delivery service, the users can be immediately presented with one or more suggested

actions in a dedicated area of user interface 112. Another surface point can be the media player experience itself. By way of example, while the user is watching a piece of content (e.g., after taking a previous action), the user can browse through alternative actions. Likewise, once the user reaches the end of a piece of content, action feed generator 108 can suggest what to do next, in other words what action to take next, and even automatically take the top suggested action on behalf of the user – creating an autoplay experience. For example, a user watching Show #1 and nearing the end of a particular episode can be suggested the actions “Continue Watching Show #1”, “Watch Another Adult Animation”, “Watch Movie #1 (movie saved for later by the user)”, and action feed generator 108 can automatically take the first action – playing the next episode of Show #1. Thus, when action feed 110 is discussed, it will be recognized that this could be a distributed list where the suggested actions in the action feed may be displayed in various different areas within the user interface 112 or may be displayed as a single list. Also, action feed generator 108 can proactively reach out to the user outside of the experience as well. For example, notifications can be sent out when an extremely relevant or time sensitive action is available. The system can be integrated with operating system (OS) and platform assistants and/or feeds.

[0033] Action feed generator 108 may personalize the action feed for every user of the video delivery service. For example, if a user #1 is using client #1, a personalized action feed #1 is displayed in a page for that user. For another user #N, client #N displays a personalized action feed #N for that user. This requires that individual user behavior be monitored and incorporated into the action feed generation process, which will be described below. Building action feed 110 may require action feed generator 108 to filter a large amount, such as tens of thousands of potential actions (at least one for each user and entity combination), down to a limited number of actions – even one. The video delivery service also has a large number of users, each of which need a personalized action feed generated. This means an automated process of generating the action feed is needed. The following will describe the different parts of the action feed generation in more detail.

[0034] Action Feed Generator

[0035] Storage Description

[0036] FIG. 2 depicts a more detailed example of action feed generator 108 according to one embodiment. Action feed generator 108 sources different information to generate actions for entities. For example, information from an action storage 202, an entity/relationship storage 204, user information storage 206, and campaign storage 208 may be used to generate actions. The information in each storage may be historical and/or real-time information. Real-time information may be information based on a current context for a user, such as the time of day, what page the user is viewing on interface 112, a selection by a user, etc. The information in each storage device will be described first and then how the information is used will be described.

[0037] Action storage 202 may store actions that can be associated with entities. The action describes a potential interaction between a user and the video delivery service. In one embodiment, actions that a user can perform on the video delivery service, such as with a remote control, may be eligible as an action in action storage 202. The actions can include both watch actions and non-watch actions. This is different from providing an icon of a video in which only one function is associated with the icon and the selection of the icon ultimately leads to watching the video. For example, the watch actions may include different types of watch actions whereas only one function is associated with the icon. The examples of actions may include categories, such as watch, save for later, follow, set a reminder, and go to a destination within the video delivery service. Each action category may have more granular actions described within it, such as “start watching”, “continue watching”, “watch latest episode”, and “watch live” are examples of watch actions in the watch category. Additionally, in the go-to actions category, the actions of “go to news hub” and “learn more” may be examples of non-watch actions. Although these actions are described, these are only examples, as other actions may be appreciated.

[0038] Examples of actions include continue watching actions, actions over followed or favorite entities, action-over-trends, recommended actions, and promoted actions. The actions may include watch and non-watch actions. For example, the non-watch actions may include follow, add an entity, or search for an entity.

[0039] The continue watching actions may generate actions for a user to resume watching a specific piece of content or continue watching a serial piece of content. For example, the action may be resume watching a movie that the user has not finished or continue watching a television show by watching the next episode.

[0040] Actions over followed or favored entities may be actions over entities that are specifically followed and/or favorited by the user. These are entities that have been specifically tracked by the user, such as the user has watched, favorited, or followed the entity. These actions may include watch the latest episode of show #1, watch a sports team vs. another team live, watch a recent interview with a person #1, watch the latest news on a current event, and so on.

[0041] The action-over-trends may generate actions that capture something trending. This may include any type of trend that the video delivery service determines by either mass engagement on the video delivery service, calling of current events on the web, or trending events that are input into the video delivery service. Some examples of actions-over-trends include watching a trending clip of a show, watch a live current event, or start watching a premiere for a television series.

[0042] Recommended actions generate actions to introduce new content to the user based on the user information, such as the user's interests, tastes, and habits. This includes start watching a television show because the user is a fan of the creator, watching the season finale of the television show because the user is into reality TV, or watching a movie because the user typically watches a movie on a Friday night.

[0043] Promoted actions generate actions that are promoted based on the video delivery service campaigns, such as campaign goals, sponsorships, or up-sales. By including promoted actions, then the entities that are included in active campaigns are considered as being surfaced as an action for a user.

[0044] Entity/relationship storage 204 may store entities and their relationships. To suggest actions for entities, relationships among entities are needed. An entity may be a video that can be played, but may be something other than a video. For example, a TV show may be an entity; an actor, director, or another person associated with the TV show may be a person entity; an adult animation may be a genre entity; a specific movie may be a movie entity; a specific user may be a user entity, and so on.

[0045] A relationship connects two entities. For example, a user #1 who watched a show #1 may have a "watched" relationship between the user #1 entity and the show #1 entity. Relationships may also have properties, such as the timestamp of when show #1 was watched

for the watched relationship. Additionally, relationships themselves may be entities, and thus there can be a relationship between entity and a relationship or between two relationships. For example, a new show can mention a guest appearance of a celebrity in a talk show, and this would be a “mention” relationship between the new show and the “appeared in” relationship between the celebrity and the talk show. An example of a relationship between two relationships would be an actor in a show being influenced by another actor in another show in an “influenced by” relationship between the two “acted in” relationships.

[0046] User information storage 206 may store a user profile, user context, user behavior, and inferences. The user information is a collection of different information about the user, which may be historical information or real-time information from a current session when the user is using the video delivery service. Each user may have personalized information stored in user information storage 206.

[0047] The user profile may include information that the video delivery service knows about the user. This information may be determined through various methods. For example, the video delivery service may receive the user profile information directly from the user in the form of questionnaires (e.g., during the registration experience), through explicit feedback from the user (e.g., whether a user likes or dislikes a video), or can be inferred from user behavior. The user profile information may be categorized into different categories, such as interests, tastes, and habits. In one example, interests may be entities that the user shows an affinity towards, such as various shows, actors, television networks, sports teams, etc. The interests may be further categorized by separating inferred interests from interests explicitly stated by the user through a mechanism in the video delivery service, such as a follow or favorite mechanism. The taste information is an abstraction over properties of entities the user is interested in, such as various comedies, movies, sports, etc. Habits are recurring behavioral patterns, such as what a user watches in the morning, weekend, at night, etc. A user profile may not change based on context. That is, the user profile information does not change whether it is 9:00 a.m. or 9:00 p.m., the weekend, or the user is in a different location.

[0048] The user context is a situation on the video delivery service associated with the user in real-time. For example, the user context may be the time of day, day of the week, the geographical location of the user, the device the user is using, the video the user is watching, etc. The user context may be determined during the session in which the user is using the video delivery service. Although the user context is described as being stored, the user

context may be automatically determined and does not need to be “stored” in user information storage 206.

[0049] The user behavior is a collection of user actions that were performed by the user. This also includes the lack of actions, such as the user decided not to start watching a video. Other user behaviors include the user started watching a video, searched for a video, navigated to a television show page, but did not watch the video. Inferences may be conclusions reached by action feed generator 108 by observing user behavior in a given context. Inferences lead to learning and these learnings may influence the user profile in interests, tastes, and habits.

[0050] Campaign storage 208 may store campaigns for the video delivery service. For example, a campaign may promote an entity, such as a media program or product. The campaign may include start and end dates, impression goals, and current progress. For example, the video delivery service may receive money for promoting a media program. The campaign may specify a number of impressions that are required for the campaign.

[0051] Candidate Action Generation

[0052] Using the above information for storage 202-208, a candidate action generator 210 generates candidate actions that are possible for a user. Candidate actions are possible actions that can be applied to entities based on a context associated with the user. For example, candidate action generator 210 generates actions that can be applied to entities at the current time based on current (or future) availability of the entities on the video delivery service. To generate the candidate actions, candidate action generator 210 may generate an entity relationship graph using the entities and relationships from entity/relationship storage 204. Then, candidate action generator 210 may perform an entity-to-content matching process to determine candidate actions for entities.

[0053] FIG. 3 shows a graph 300 of entities and the relationships between the entities according to one embodiment. In one embodiment, graph 300 may represent any content that may be shown regardless of whether or not that content is available on the video delivery service. Graph 300 includes entities that may not be available on the video delivery service because there may be relationships between entities currently on the video delivery service

and entities not currently on the service that may be drawn, or there may be relationships to future availability.

[0054] In one embodiment, every entity in graph 300 is unique. For example, at 302, a show #1 is a unique entity, and at 304, person #1 is a unique entity. Graph 300 may be a visualization of a subset of entities and relationships that provides the available possible entities and relationships. For example, a show #1 at 302 may have a relationship 306 of “has an episode from” season 10, episode 8 (S10, E8) at 303. Further, show #1 may have a relationship 307 of “created by” a person #1 that is shown at 304. Then, the person #1 may also have a relationship 308 of “appeared on” a television episode of season 3, episode 57 (S3, E57) of a television show #2 at 310. Television show #2 at 312 also has a relationship 313 of “has episode from” the same episode of S3, E57. Further, a movie #1 at 314 may have a relationship 316 with person #1 in that person #1 created movie #1.

[0055] Candidate action generator 210 uses graph 300 to understand what possible entities/relationships could occur whether or not the entity may be available on the video delivery service. Then, candidate action generator 210 uses an entity-to-content matching process that associates certain types of entities in the entity-relationship graph with actual content that can be requested by a user using the video delivery service. For a video title, the video delivery service may have several sources for viewing the video, such as on-demand and linear. Further, different sources might require different entitlements, such as different subscription levels for the video delivery service. Candidate action generator 210 uses the entity-to-content matching process to determine which entities are available for a user to watch. The availability may be availability at a particular point in time, past availability, or future availability.

[0056] FIG. 4 depicts an example of an entity-to-content matching process 400 according to one embodiment. Entity-relationship graph 300 is shown with a content pipeline 402 for the video delivery service. Content pipeline 402 may describe the content that is available at the current time or in the future on the video delivery service. For example, the entity of show #1, which has an episode of season 10, episode 8 is available in content pipeline 402 in different forms. For example, at 406, the episode may be available on-demand. This means a user can currently request the episode at any time. Further, at 408 and 410, the episode may be available via linear delivery, such as the episode is available live on a Channel #1 at a first time #1 at 408 and available live on a Channel #2 at a second time #2 at 410. Candidate

action generator 210 determines the entities at 406, 408, and 410 when performing the entity-to-content matching process 400.

[0057] Entity-to-content matching process 400 uses current availability during the candidate generation, but future availability is also considered. That is, although some entities may not be currently available, it is possible that these entities may be available in the future on the video delivery service. For example, although television shows, people, etc. are entities that do not have corresponding videos (unlike television episodes or movies), candidate action generator 210 can traverse the relationships in graph 300 and generate actions over these entities if they have relationships to watchable entities that are linked to the available video. For example, if a person #1 appears in a television episode and the video for that television episode is available, the action can be over person #1 and could suggest the video with person #1 in it.

[0058] Referring back to FIG. 2, after performing entity-to-content matching process 400, candidate action generator 210 can generate candidate actions for the available entities. Candidate action generator 210 can use the available entities determined by the entity-to-content matching process 400 in addition to the user information from user information storage 206 to determine candidate actions for entities. For example, when a user logs on to the video delivery service, or at other points while the user is using the video delivery service, candidate action generator 210 may generate candidate actions based on the user profile and also the context of the user. The candidate actions may be generated at other times also, such as when the user navigates to a different page of interface 112. The candidate actions may not all be seen by the user, but candidate action generator 210 may determine they could qualify as potential actions to surface to the user.

[0059] In one embodiment, candidate action generator 210 can determine entities a user may be interested in and also generate actions for the entities based on what content is available via entity-to-content matching process 400. In one example, candidate action generator 210 may use a machine-learning predictor to generate actions for the entities. In another example, rules may also be used in addition to or in lieu of the machine-learning predictor. In one embodiment, candidate action generator 210 can generate the possible actions for entities based on the user watch behavior on the video delivery service and a real-time user context. For example, candidate action generator 210 may predict whether a user is interested in an entity the user has interacted with before, predict whether a user will take a certain action on

an interested entity the user has interacted with before, and/or predict whether a user is interested in a new entity when the user navigates to a page on interface 112. In other embodiments, candidate action generator 210 may determine whether there is an available action for an entity based on the entity's status without taking into account specific user behavior. For example, if a trailer is released for a movie, then the entity/action pair might be watch the trailer for the movie.

[0060] In the rules-based predictor system, candidate action generator 210 may use user behavior to predict a user is interested in an entity. For example, candidate action generator 210 predicts a user is interested in a show if the user has completed one episode. The entity-to-content matching process indicates which shows are offered on the video delivery service and also metadata for the shows, such as the actors in the shows, number of episodes, etc. Then, candidate action generator 210 may generate a set of actions for the show, such as watch the latest episode when a new episode (assuming the user is caught up with all released episodes) or watch the next episode if there is a second released episode for the show.

[0061] In the machine-learning predictor system, candidate action generator 210 may predict whether a user is interested in different types of entities, such as shows, movies, genres, actors, directors, networks, etc. For example, candidate action generator 210 may use the user behavior and information for the available entities to predict whether a user would re-engage with these entities from user behavior on the video delivery service, such as video views and also metadata for the entities. Also, with the help of the entity/relationship graph, connections between a user and an entity that have not been engaged by the user can be drawn. Based on connections, candidate action generator 210 may recommend new entities and appropriate associated actions (Follow "Actor A", Follow "Election 2016").

[0062] The machine-learning predictor may first be trained using historical user behavior (e.g., user actions, such as browse behavior, watch behavior, and search behavior) on the video delivery service and the information for the available entities. This trains a model to predict whether a user would re-engage with an entity based on the user's short-term or long-term engagement with related entities. These entities may include shows, which may be television shows, manually-followed entities, and other entities. If the predicted probability of re-engaging with an entity is higher than a predefined threshold, then candidate action generator 210 may consider this entity as an entity the user is interested in.

[0063] In one embodiment, candidate action generator 210 selects a first set of entities that the user has engaged within a time period (e.g. the last 3 months). Candidate action generator 210 may also determine whether the probability of the user re-engaging with the entity is above a threshold. Also, candidate action generator 210 selects a second set of entities the user has not engaged with (e.g., with the last 3 months or ever), but are highly related to what the has engaged before. Candidate action generator 210 determines the second set of entities dynamically, based on a context for the user, such as the time of the day, geo-location, if the user has just finished watching a show, or other contextual information.

[0064] After selecting the entities, candidate action generator 210 generates action candidates for the predicted entities. In one embodiment, all possible actions are applied, such as the actions are ones that can be applied to the entities at the current time or a future time. In another embodiment, candidate action generator 210 predicts what actions the user might use to re-engage with the entity. For example, watch actions may be associated with the show that has available episodes to watch for a specific user. For manually-followed entities, the action may be watch the entity from the beginning. Other actions may be recommended for entities that the user may be interested in, such as follow an actor that is associated with a show the user watched. Candidate action generator 210 may then output the predicted possible actions for the entities. It is noted that entities may have multiple actions that could be applied. The actions output by candidate action generator 210 are received at action ranker 212 for ranking.

[0065] Action Ranking

[0066] After generating candidate actions over entities, an action ranker 212 ranks the candidate actions against each other. Action ranker 212 may rank the actions themselves against each other instead of ranking content against content. That is, the action of watching a television episode may be ranked against the action of watching a movie, instead of ranking the television show vs. the movie. The ranking of actions allows action ranker 212 to factor in user context effectively and capture habits. This is different from ranking a movie against a television show, and requires a different type of predictor to perform the ranking. For example, action ranker 212 may recommend an action such that the probability of the user selecting the action in the current context is maximized. Further, action ranker 212 may

recommend an action such that the downstream value of the user taking the action is maximized. The downstream value may be some value of the user taking this action at a later time. This value may be quantified on the video delivery service. For example, not only does the video delivery service want to provide an engaging experience to a user currently using the video delivery service, the video delivery service would like to provide a content strategy for promoting its content organically in the future. For example, an action the user might most likely take is watching a commonly available television show. However, an exclusively available critically acclaimed show might be slightly lower in the probability of the user currently taking that action, but would have a much higher downstream value in that the user might continue to view this more valuable show later. In this case, the critically acclaimed show might be ranked higher. The actions also go beyond suggesting that a user watch a video as non-watching actions are included in the ranking process with watch actions.

[0067] FIG. 5A depicts an example of the ranking process 500 used by action ranker 212 according to one embodiment. Action ranker 212 may include multiple rankers that can rank entities. The process of ranking may take into account information from entity/relationship storage 204, user information storage 206, and campaign storage 208. This information is used to generate candidate actions for entities as was described above. The candidate generation process will not be described again, but the specific information used to generate candidates for each ranker is described. For example, user information storage 206 may include user behavior, manually-added entities, and engaged entities shown at 502, 504, and 506, respectively. As discussed above, the user behaviors in the past include browse, search, watch, follow, and save behaviors. The user behavior is used to determine manually-added entities 504 and engaged entities 506. Manually-added entities 504 are those entities that a user has manually added, followed, or saved. Manually-added entities 504 may be explicitly added by the user. Engaged entities 506 may be those entities that a user has watched, searched for, or browsed over. The engaged entities may be implicitly expressed by the user.

[0068] Content relevance at 508 stores relevance between two pieces of content or two entities. This is used to determine if a user is interested in one media program, the user might be interested in another media program.

[0069] At 510, the content availability is the availability of content including when the content premieres, expires, and what time slot that live content is going to be provided. The

content availability is determined from the entity/relationship storage 204. The content availability may be provided for all combinations of geo-location and video delivery service package types. That is, some content may be available in certain locations and some content may be available in different packages offered by the video delivery service.

[0070] At 512, the campaign information may be retrieved from campaign storage 208. The campaign information may be information on campaigns that are being run relating to a certain entity. The campaign information includes the start and end dates, impression goals and current progress, and also actions that can be taken for the entity that is the subject for the campaign.

[0071] Action ranker 212 receives the candidate actions for the information retrieved from storage 502-512 at the different rankers. The rankers are configured to receive actions for entities and automatically generate a probability a user would select each action. Because of the nature of the user using the video delivery service and navigating to different areas of interface 112, the rankers need to be used to automatically generate the rankings so the actions can be surfaced to a user in real-time. An example of a ranker that can be used will be described in more detail below.

[0072] A re-engagement ranker 514 receives the manually-added entities 504 and engaged entities 506. Then, re-engagement ranker 514 ranks actions for the entities for which the user has knowledge. This ranks the probability that the user would select the action for an entity that the user has explicitly added or taken action on before.

[0073] A discovery ranker 516 receives the engaged entities and content relevance information. Discover ranker 516 ranks actions for the entities for which the user does not have any knowledge. For example, the engaged entities are used to determine relevant entities that are related to the engaged entities. Discovery ranker 516 discovers the entities based on entity/relationship graph 300.

[0074] A campaign urgency ranker 518 receives the campaign information and can rank the actions for the campaigns. In one embodiment, campaign urgency ranker 518 may rank an action for a campaign based on an urgency score for the campaign. The urgency may be how important it is to receive an impression for the entity at the particular point in time. The higher the urgency score is, the more urgent the campaign is. The score may be computed

based on the campaign goals and the current progress. For example, the campaign goal may be 1 million impressions in 1 week, and the current progress may be 100,000 impressions after 3 days. This may make the importance of receiving an impression for this campaign more urgent than if 800,000 impressions had already been received.

[0075] The output of re-engagement ranker 514 and discovery ranker 516 is input into a user entity relevance ranker 520. Additionally, user entity relevance ranker 520 receives the content availability. Then, user entity relevance ranker 520 generates a unified ranking of the inputs with each of the entities' actions being ranked based on a probability the user would select each action. That is, the explicit actions and implicit actions are ranked against each other. In one embodiment, the explicit actions and implicit actions can be combined using different methods. For example, re-engagement ranker 514 and discovery ranker 516 output probabilities within the same range (e.g., between a range of [0,1]) and thus the actions are comparable. In other examples, the probabilities from one of the rankers may be normalized such that the probabilities for both rankers are within the same range after normalization. Once the probabilities are in the same range, user entity relevance ranker 520 can unify the ranking in a descending order according to the probabilities output by the previous rankers (e.g., re-engagement ranker 514 and discovery ranker 516). Another way of unifying could use a pre-specified ratio of re-engagement vs. discovery to merge the ranked list from the previous rankers according to the ratio. For example, with a 3:1 ratio, user entity relevance ranker 520 selects the three highest ranked actions from re-engagement ranker 514, and then selects the single (1) highest ranked action from discovery ranker 516, and then selects the next three highest ranked actions from re-engagement ranker 514, so on. Other considerations may also be used to unify the ranking, such as user entity relevance ranker 520 can promote an action based on some context, such as special treatment can be made when user entity relevance ranker 520 looks at the user's behavior, e.g. when the user has finished one show, then the ratio of discovery item can be increased.

[0076] A campaign ranker 522 receives the campaign urgency ranking along with the content relevance and content availability. Then, campaign ranker 522 may rank all campaigns with relevance between the content and the campaign with the urgency considered. For example, the relevance of the content for the campaign being promoted along with the content availability is used to rank actions for all the campaigns.

[0077] Then, a mixed ranker 524 receives the user entity relevance ranking and the campaign ranking. Mixed ranker 524 can rank a final list of recommended entities with associated actions from campaign ranker 522 and user entity relevance ranker 520. Mixed ranker 524 may unify the rankings from campaign ranker 522 and user entity relevance ranker 520 in a similar way as described with respect to user-entity relevance ranker 520. For example, campaign ranker 522 and user entity relevance ranker 520 may output probabilities for actions within the same range of [0,1] that are ranked together, or use any of the other methods described.

[0078] Each ranker discussed above may use a predictor to rank actions. Each different predictor may rank a different set of actions, but may use similar a ranking process. The ranking process will be described with respect to action ranker 212, but may be used by each ranker described in FIG. 5A.

[0079] In one embodiment, action ranker 212 may use a machine-learning predictor. Also, a rules-based predictor may also be used in addition to or in lieu of the machine-learning algorithm. For the rules-based predictor, action ranker 212 may rank the actions using historical user behavior on the video delivery service. The historical user behavior data may be associated with the user, and/or be associated with other users on the video delivery service. That is, data may be aggregated for many users using the video delivery service and used to predict a ranking for the actions for entities. Action ranker 212 then applies rules to the real-time context to rank the actions for the entities.

[0080] For the machine learning predictor, action ranker 212 may be trained with previous user behaviors from the past. Each behavior can be interpreted as an action on a specific entity. All the actions, associated with a context when the user took the action, form the positive samples. Negative samples are those at the same context that were eligible to be taken by the user, but were not taken by the user. The training process tries to find the predictor with the best discriminative power to differentiate positive samples from negative ones, using a prediction process, such as multi-layer neural networks, decision tree or logistic regression. As a result, candidate action generator 210 outputs a value for a specific action, which is the predicted probability that the action would be taken by the user. There might be more than one action eligible for a user to take on an entity, but action ranker 212 may only output the action with highest predicted probability for each entity. However, action ranker 212 can also output multiple actions for entity.

[0081] In one example, the trained predictor includes a function $f(u, e, a, c)$, which outputs a value between a range (e.g., $[0,1]$), where u is the user (including information about the user's past behaviors up to a certain time point), e is the entity, a is an action, and c is the real-time context (time of day, geo-location (e.g., location information), device being used, seasonality, day of the week, current status of the user (e.g., just finished a show episode), etc.). Action ranker 212 retrieves the information that is needed or receives the information dynamically in real-time. The lower the value output by the predictor, the less likely the user is predicted to take the action, and the higher the value, the more likely the user is predicted to take the action.

[0082] When given a set of pairs of form (e, a) for a certain user u in the context c , the predictor evaluates all values of $f(u, e, a, c)$. Then, action ranker 212 ranks all the (e, a) pairs according to the evaluated values. For each entity e , only one (e, a) pair may be retained in the final results, namely, only the most probable action for an entity is retained. However, multiple actions may be used. In addition to relevance, a user's negative feedback may also be used by action ranker 212 to adjust the ranking.

[0083] In one embodiment, the inputs for the entity, user, or context to the predictor may be based on any combination of:

Attributes of the entity (e).

Availability change for the entity (e).

The user's context (c).

The user's affinity to the entity (u).

The user's recent behaviors (u).

The user's habits on the video delivery service (u).

[0084] The attributes of the entity may include information about a video, such as when a show is releasing episodes, when the last episode was released, what episode the user last watched, etc. The availability change may be if any availability on the video delivery service might change, such as the episode or movie may be scheduled to be removed from the video delivery service. The user's context is based on various characteristics for the user, such as

the device being used to access the video delivery service, the time of the day, the day of the week, etc. The user's affinity to the entity may be how much the user likes the entity, such as whether the user has watched a related video, liked the director or other director's movies, liked the show, etc. The user's recent behaviors may be what actions a user has recently taken on the video delivery service, such as manually added/followed a video, watched a video, etc. The user's habits may be actions a user typically takes on the video delivery service.

[0085] Action ranker 212 may generate probabilities that a user will take an action for an entity by ranking the actions for entities against each other. For an entity that the user has not previously performed an action with, candidate action generator 210 may infer the probability by the relevance between this entity/action and entities the user has taken actions on. For an entity the user has performed actions on before, action ranker 212 may predict the probability based on the user's watch behavior, watch patterns, search behaviors, and availability of the entity.

[0086] Action ranker 212 may also use the effective time to determine how likely the user will take the action, such as it may be more likely that a user will re-engage with an entity to perform an action if the user more recently performed that action on the same entity. The effective time may be dependent on the type of entity. For example, effective time may be defined as a function of last engaged time or availability change time (new arrival, expiring soon, season premiere, and a manually-saved time or followed time for the user). The effective time may also be based on a number of unseen episodes for the specific user. In one example, if the user's favorite show or celebrity has new arrivals in the video delivery service (e.g., a new episode or movie), the user saved or watched half-way a show/movie recently, the user has been binge watching a library show, action ranker 212 may generate a high probability for the entity/action pair.

[0087] Action ranker 212 may use the context to determine probabilities that the user will take the action on the entity, such as the user may use a certain device to access the service at a certain time. The context may be the user's location, time, device that the user is currently using, and the subscription package for the video delivery service.

[0088] Conventionally, when ranking content against each other, the system is concerned with whether the user engages with the content. However, ranking actions takes the

prediction one step further, not only is action ranker 212 concerned about the content the user would engage with, but also how the user would engage with the content. Therefore, in one embodiment, the predictor can explicitly tell the user the suggested way of engagement with the entity.

[0089] The output of action ranker 212 may be probabilities for entities and each associated action pairs. FIG. 5B depicts a table 550 of entities and actions that have been ranked according to one embodiment. Table 550 may be stored in a database and accessed by action feed generator 108 to generate a list of actions for entities for the action feed. A column 552 lists the entities that were selected by candidate action generator 210 and column 554 lists the actions that were selected for the entities by candidate action generator 210. Column 556 lists the probabilities generated by action ranker 212. Each show/action entity may be a row in the table although this is not shown graphically. The probabilities indicate a probability the user would take an action over another action. The probabilities may be based on all actions in the tables. For example, for the action Watch next episode for Show #1, it is 90% probable the user would take this action over all other actions (or a group of actions) in the table.

[0090] Action Feed Filtering

[0091] The output of mixed ranker 524 is an action feed that includes a ranked list of actions on entities. This may be considered an action feed 110 that can be used by the video delivery service. For example, action feed 110 may be surfaced to various interface positions at a client 104. Additionally, action feed 110 may be filtered to provide a drill-down into different categories of actions. For example, the action feed may be provided on a homepage for a user when the user logs into the video delivery service. Also, action feed 110 may be filtered to only show actions for entities that are new to the user; that is, entities that a user has explicitly watched or added before and does not include campaigns.

[0092] FIG. 6A depicts a more detailed example of action feed generator 108 for displaying action feed 110 according to one embodiment. An action feed trigger 601 receives a real-time indication that the action feed should be displayed on interface 112. For example, the

trigger may be that a user has logged onto the video delivery service. Other triggers include the user browsing to different pages, hubs, or sections of user interface 112.

[0093] A filter 602 receives action feed 110 from mixed ranker 524 when the trigger is detected. Filter 602 may also receive real-time information for a user, such as a user's context. For example, the context may include the page of an interface the user is currently viewing, current actions the user is performing, or other user information. The user context may be received from the trigger, from client 104, or derived (e.g., a time of day).

[0094] Filter 602 may then filter the action feed to select a portion of the actions based on the user's context. In one embodiment, filtering may remove some actions that are higher ranked than others based on the filter criteria. Although this type of filtering is described in other embodiments, the action feed may select the top X actions from action feed 110. In one example, only one action is displayed per entity. In this example, X is a number equal to or less than the number of entities in the action feed. In another example, multiple actions may be displayed for one or more entities. In the second example, X is a number equal to or less than the number of actions in the action feed.

[0095] Filter 602 may then generate a table that includes the selected action feed entities and actions. FIG. 6B depicts an example of a table 650 that is used to store the selected actions according to one embodiment. As shown, table 650 includes the selected entities in a column 652 and also the actions in a column 654. Because different actions are possible, table 650 includes a link to the selected action/entity combination in a column 656. For example, different actions may point to different episodes of a show. Depending on the action, the link points to the correct entity, such as the latest episode for a caught up show or the first episode for a library show.

[0096] In the example in table 650, the actions are listed from highest probability to the lowest probability. For show #1, a link #1 includes information such that the next episode for show #1 that a user still has not watched is played. This link may be hardcoded to a specific episode or the next episode may be determined dynamically when link #1 is selected. For show #2, a link #2 causes show #2 to be followed. For the Drama genre, a link #3 causes the newest release in the drama genre to be played. For Actor #1, a link #4 causes the latest update (e.g., an episode including the Actor #1) to be played.

[0097] Referring back to FIG. 6A, whatever actions that are selected by filter 602 are then provided to display process 604. Display process 604 can then dynamically display the action feed whether it is filtered or not on interface 112 of client 104. Display process 604 may dynamically display action feed 110 while the user navigates to different areas of interface 112 and different pages are displayed.

[0098] In one example, action feed 110 can be displayed on a homepage for the user in interface 112. The homepage may be the first page of an application that a user sees when the user logs into the video delivery service. In one embodiment, this action feed is displayed in an area, such as the masthead or a prominent location on interface 112. Additionally, action feed 110 may also be provided on different trays, hubs, or pages of interface 112. For example, different filtered action feeds can be provided in different pages. Examples of pages include a new-for-you category that shows only actions that are new to the user. Display process 604 may also display slices of the action feed as the user browses for something specific. For example, if the user navigates to a movies page on interface 112, then the action feed can be filtered down to only movie actions. Display process 604 can display these movie actions on the current page of the user.

[0099] Display process 604 may also display an action feed for continuous play. For example, when a user reaches the end of playback of a video, action feed trigger 601 receives the notification that playback is ending, and display process 604 can select the next best action for the user to display on interface 112. This action could keep the user watching without any explicit effort by the user and also suggest other alternatives or pivots for the user to view. The pivots can be categorized similarly to the approach that generated the actions for the homepage. When performing continuous play, the prior action is considered when generating and ranking the actions. If the user has been viewing a lot of episodes in a row of a television show, then the action should be continue watching that television show as a top-ranked suggested action.

[00100] Notifications may also be used when a top action is very relevant and potentially time-sensitive, such as watching a live sporting event. Display process 604 may then raise a notification to the user in interface 112. This may be in interface 112 and may also be through the operating system such that this is a global notification for the user.

[00101] Interface and Action Feed Example

[00102] FIG. 7 depicts an example of action feed 110-1 being displayed on a user interface 112 of client 104 according to one embodiment. Although this user interface is shown, it will be understood that users may navigate to other areas of an application being displayed in the user interface to display different versions of the action feed. In this version, interface 112 may display different categories 704-1 - 704-4 of videos from the video delivery service's video library. For example, different categories may include a lineup that shows top suggested actions for the user, a television category, a movies category, a sports category, and a networks category. Each category includes its own action feed. Although these categories are listed, other categories may also be appreciated.

[00103] In one embodiment, a selection of category #1 is received using a selector 706. The selection causes display of action feed 110-1, which includes entities #1 - #4 702-1 - 702-4 arranged in a predicted order for actions 708-1 - 708-4 for the entities. For example, at are shown in an order from top to bottom. For example, entity 702-1 has the highest probability the user will perform the associated action #1 708-1, entity 702-2 has the next highest probability the user will perform the associated action #1 708-2, etc. Also, other entities with associated actions may also be included in action feed 110 after action #4 708-4, but are not shown.

[00104] When entities are shows, the shows may include a different number of available or unwatched episodes. The associated action may suggest an episode that has the highest probability the user might want to watch next. Also, live videos may be blended with on demand videos in the same action feed 110-1. The ranking of actions allows the system to rank the live videos with the on demand videos and provide a unified interface with both types of videos.

[00105] In one embodiment, action feed generator 108 uses table 550 shown in FIG. 5B to dynamically generate action feed 110-1 in interface 112. For example, action feed generator 108 determines how many spots in the category are available for action feed 110-1. The number of spots may depend on the category or on the number of entities and associated probabilities in table 550. For example, action feed generator 108 may select each unique entity that is included in table 550. Then, action feed generator 108 selects the action for each unique entity that has the highest probability. Action feed generator 108 may also use

other methods to select the entities, such as selecting X number of entities that have actions above a threshold.

[00106] Once selecting the entities, action feed generator 108 uses the links to the content in table 650 of FIG. 6B to generate action feed 110. For example, the links are sent to client 104 and when an entity is selected, client 104 uses the link to request the associated entity with the action. For example, the link not only specifies the entity, but also the action. Link #1, #3, #4 will direct a user to a page where a video is being played. Link #2 will direct a user to a page describing details of the show and a message indicating the future updates of the show will be displayed in the user's action feed 110.

[00107] Method Flows

[00108] FIG. 8 depicts a simplified flowchart 800 of a method for generating an action feed 110 according to one embodiment. At 802, action feed generator 108 sends content to users. At 804, action feed generator 108 stores user behavior on the video delivery service based on the sending of the content to user. The user behavior includes actions taken by the users (or not taken) on the video delivery service.

[00109] Then, for each user, at 806, action feed generator 108 inputs the user behavior, a user context, entity/relationship information, and campaign information into candidate action generator 210 to generate a set of candidate actions to be performed on a set of entities found on the video delivery service. At 808, action feed generator 108 inputs the set of candidate actions to be performed on the set of entities, the user context, and the user behavior into action ranker 212 to generate probabilities for the set of actions to be performed on the set of entities. As discussed above, the probability for an action indicates the probability the user would select that action when compared against other actions in the set of actions to be performed on entities. At 810, information for the probabilities for the set of actions is stored in a table.

[00110] After generating the ranked action feed, FIG. 9 depicts a simplified flowchart 900 of a method for displaying action feed 110 according to one embodiment. At 902, action feed generator 108 receives a trigger for the action feed. The trigger indicates the action feed should be generated.

[00111] At 904, action feed generator 108 determines a user context. The user context may be based on real-time information for the user. At 906, action feed generator 108 selects actions for an action feed 110 based on the ranking. For example, depending on the context, action feed generator 108 applies different filters to action feed 110. At 908, action feed generator 108 outputs action feed 110 to the users.

[00112] FIG. 10 depicts a simplified flowchart 1000 of a method for responding to a selection of an action in action feed 110 according to one embodiment. At 1002, action feed generator 108 receives a selection for an action on an entity in action feed 110. At 1004, action feed generator 108 determines a link to the action. For example, client device 104 may send the link to the action in a request. The link may have been associated with the icon on the action feed that was selected. In other examples, action feed generator 108 determines the link once the request for the action is received for table 650 in FIG. 6B. At 1006, action feed generator 108 performs the action. For example, if the action is watch the next episode of a show, then action feed generator 108 causes the next episode of the show to be played on video delivery system 106.

[00113] Feedback

[00114] When actions are provided to a user, feedback may be received and used to improve the generation of future action feeds. Action feed generator 108 receives the actions that a user takes. The feedback may also be negative feedback in that the user did not take an action that was suggested. Positive feedback may be when the user actually takes an action. The positive feedback may be binary in that the user took an action or did not take an action, or it can be weighted by downstream value/intensity of engagement. For example, the user may take an action to start viewing a television show but then stop viewing the television show a few minutes later in contrast to the user viewing every single episode of the season. Also, the negative feedback is when the user ignores a recommended action or through explicit feedback such as the user selects a dismiss or not interested input. The feedback information is then provided in the user behavior database for future use in generating the action feed.

[00115] Conclusion

[00116] Accordingly, the action feed can be used to suggest actions for a user. This is different from suggesting just videos. The actions take the recommendations a step further to suggest what the user could do with the entities. Additionally, the ranking process ranks the actions against each other. This allows the recommendations to use user behavior with more granularity to provide more accurate recommendations.

[00117] Further, the actions allow the entities to be ranked from multiple sources, such as linear video in addition to on-demand video. Additionally, campaigns can also be integrated into the ranking or the action feed.

[00118] The detection of patterns in historical user behavior can then be used to more highly rank actions that fit the user's behavior. As a result, fewer actions over entities can be recommended, as opposed to the entity itself that the video delivery service knows is related to a video that the user has previously watched. This also allows context to be provided, such as if the user watches news in the morning, then actions over the news content can be provided to the user.

[00119] System Overview

[00120] Features and aspects as disclosed herein may be implemented in conjunction with a video streaming system 1100 in communication with multiple client devices via one or more communication networks as shown in FIG 11. Aspects of the video streaming system 1100 are described merely to provide an example of an application for enabling distribution and delivery of content prepared according to the present disclosure. It should be appreciated that the present technology is not limited to streaming video applications, and may be adapted for other applications and delivery mechanisms.

[00121] In one embodiment, a media program provider may include a library of media programs. For example, the media programs may be aggregated and provided through a site (e.g., Website), application, or browser. A user can access the media program provider's site or application and request media programs. The user may be limited to requesting only media programs offered by the media program provider.

[00122] In system 1100, video data may be obtained from one or more sources for example, from a video source 1110, for use as input to a video content server 1102. The input video data may comprise raw or edited frame-based video data in any suitable digital format, for example, Moving Pictures Experts Group (MPEG)-1, MPEG-2, MPEG-4, VC-1, H.264 / Advanced Video Coding (AVC), High Efficiency Video Coding (HEVC), or other format. In an alternative, a video may be provided in a non-digital format and converted to digital format using a scanner and/or transcoder. The input video data may comprise video clips or programs of various types, for example, television episodes, motion pictures, and other content produced as primary content of interest to consumers. The video data may also include audio or only audio may be used.

[00123] The video streaming system 1100 may include one or more computer servers or modules 1102, 1109, and/or 1107 distributed over one or more computers. Each server 1102, 1109, 1107 may include, or may be operatively coupled to, one or more data stores 1104, for example databases, indexes, files, or other data structures. A video content server 1102 may access a data store (not shown) of various video segments. The video content server 1102 may serve the video segments as directed by a user interface controller communicating with a client device. As used herein, a video segment refers to a definite portion of frame-based video data, such as may be used in a streaming video session to view a television episode, motion picture, recorded live performance, or other video content.

[00124] In some embodiments, a video advertising server 1104 may access a data store of relatively short videos (e.g., 10 second, 30 second, or 60 second video advertisements) configured as advertising for a particular advertiser or message. The advertising may be provided for an advertiser in exchange for payment of some kind, or may comprise a promotional message for the system 1100, a public service message, or some other information. The video advertising server 1104 may serve the video advertising segments as directed by a user interface controller (not shown).

[00125] The video streaming system 1100 also may include action feed generator 108.

[00126] The video streaming system 1100 may further include an integration and streaming component 1107 that integrates video content and video advertising into a streaming video segment. For example, streaming component 1107 may be a content server or streaming media server. A controller (not shown) may determine the selection or

configuration of advertising in the streaming video based on any suitable algorithm or process. The video streaming system 1100 may include other modules or units not depicted in Fig. 11, for example administrative servers, commerce servers, network infrastructure, advertising selection engines, and so forth.

[00127] The video streaming system 1100 may connect to a data communication network 1112. A data communication network 1112 may comprise a local area network (LAN), a wide area network (WAN), for example, the Internet, a telephone network, a wireless cellular telecommunications network (WCS) 1114, or some combination of these or similar networks.

[00128] One or more client devices 1120 may be in communication with the video streaming system 1100, via the data communication network 1112 and/or other network 1114. Such client devices may include, for example, one or more laptop computers 1120-1, desktop computers 1120-2, “smart” mobile phones 1120-3, tablet devices 1120-4, network-enabled televisions 1120-5, or combinations thereof, via a router 1118 for a LAN, via a base station 1117 for a wireless telephony network 1114, or via some other connection. In operation, such client devices 1120 may send and receive data or instructions to the system 1100, in response to user input received from user input devices or other input. In response, the system 1100 may serve video segments and metadata from the data store 1104 responsive to selection of media programs to the client devices 1120. Client devices 1120 may output the video content from the streaming video segment in a media player using a display screen, projector, or other video output device, and receive user input for interacting with the video content.

[00129] Distribution of audio-video data may be implemented from streaming component 1107 to remote client devices over computer networks, telecommunications networks, and combinations of such networks, using various methods, for example streaming. In streaming, a content server streams audio-video data continuously to a media player component operating at least partly on the client device, which may play the audio-video data concurrently with receiving the streaming data from the server. Although streaming is discussed, other methods of delivery may be used. The media player component may initiate play of the video data immediately after receiving an initial portion of the data from the content provider. Traditional streaming techniques use a single provider delivering a stream of data to a set of end users. High bandwidths and processing power may be required to

deliver a single stream to a large audience, and the required bandwidth of the provider may increase as the number of end users increases.

[00130] Streaming media can be delivered on-demand or live. Streaming enables immediate playback at any point within the file. End-users may skip through the media file to start playback or change playback to any point in the media file. Hence, the end-user does not need to wait for the file to progressively download. Typically, streaming media is delivered from a few dedicated servers having high bandwidth capabilities via a specialized device that accepts requests for video files, and with information about the format, bandwidth and structure of those files, delivers just the amount of data necessary to play the video, at the rate needed to play it. Streaming media servers may also account for the transmission bandwidth and capabilities of the media player on the destination client. Streaming component 1107 may communicate with client device 1120 using control messages and data messages to adjust to changing network conditions as the video is played. These control messages can include commands for enabling control functions such as fast forward, fast reverse, pausing, or seeking to a particular part of the file at the client.

[00131] Since streaming component 1107 transmits video data only as needed and at the rate that is needed, precise control over the number of streams served can be maintained. The user will not be able to view high data rate videos over a lower data rate transmission medium. However, streaming media servers (1) provide users random access to the video file, (2) allow monitoring of who is viewing what video programs and how long they are watched, (3) use transmission bandwidth more efficiently, since only the amount of data required to support the viewing experience is transmitted, and (4) the video file is not stored in the user's computer, but discarded by the media player, thus allowing more control over the content.

[00132] Streaming component 1107 may use TCP-based protocols, such as HTTP and Real-time Messaging Protocol (RTMP). Streaming component 1107 can also deliver live webcasts and can multicast, which allows more than one client to tune into a single stream, thus saving bandwidth. Streaming media players may not rely on buffering the whole video to provide random access to any point in the media program. Instead, this is accomplished through the use of control messages transmitted from the media player to the streaming media server. Another protocol used for streaming is hypertext transfer protocol (HTTP) live streaming (HLS). The HLS protocol delivers video over HTTP via a playlist of small

segments that are made available in a variety of bitrates typically from one or more content delivery networks (CDNs). This allows a media player to switch both bitrates and content sources on a segment-by-segment basis. The switching helps compensate for network bandwidth variances and also infrastructure failures that may occur during playback of the video.

[00133] The delivery of video content by streaming may be accomplished under a variety of models. In one model, the user pays for the viewing of video programs, for example, using a fee for access to the library of media programs or a portion of restricted media programs, or using a pay-per-view service. In another model widely adopted by broadcast television shortly after its inception, sponsors pay for the presentation of the media program in exchange for the right to present advertisements during or adjacent to the presentation of the program. In some models, advertisements are inserted at predetermined times in a video program, which times may be referred to as “ad slots” or “ad breaks.” With streaming video, the media player may be configured so that the client device cannot play the video without also playing predetermined advertisements during the designated ad slots.

[00134] Referring to FIG. 12, a diagrammatic view of an apparatus 1200 for viewing video content and advertisements is illustrated. In selected embodiments, the apparatus 1200 may include a processor (CPU) 1202 operatively coupled to a processor memory 1204, which holds binary-coded functional modules for execution by the processor 1202. Such functional modules may include an operating system 1206 for handling system functions such as input/output and memory access, a browser 1208 to display web pages, and media player 125 for playing video. The modules may further include modules to generate interface 112 and action feed 110. The memory 1204 may hold additional modules not shown in FIG. 12, for example modules for performing other operations described elsewhere herein.

[00135] A bus 1214 or other communication component may support communication of information within the apparatus 1200. The processor 1202 may be a specialized or dedicated microprocessor configured to perform particular tasks in accordance with the features and aspects disclosed herein by executing machine-readable software code defining the particular tasks. Processor memory 1204 (e.g., random access memory (RAM) or other dynamic storage device) may be connected to the bus 1214 or directly to the processor 1202, and store information and instructions to be executed by a processor 1202. The memory

1204 may also store temporary variables or other intermediate information during execution of such instructions.

[00136] A computer-readable medium in a storage device 1224 may be connected to the bus 1214 and store static information and instructions for the processor 1202; for example, the storage device (CRM) 1224 may store the modules 1206, 1208, 1210 and 1212 when the apparatus 1200 is powered off, from which the modules may be loaded into the processor memory 1204 when the apparatus 1200 is powered up. The storage device 1224 may include a non-transitory computer-readable storage medium holding information, instructions, or some combination thereof, for example instructions that when executed by the processor 1202, cause the apparatus 1200 to be configured to perform one or more operations of a method as described herein.

[00137] A communication interface 1216 may also be connected to the bus 1214. The communication interface 1216 may provide or support two-way data communication between the apparatus 1200 and one or more external devices, e.g., the streaming system 400, optionally via a router/modem 1226 and a wired or wireless connection. In the alternative, or in addition, the apparatus 1200 may include a transceiver 1218 connected to an antenna 1229, through which the apparatus 1200 may communicate wirelessly with a base station for a wireless communication system or with the router/modem 1226. In the alternative, the apparatus 1200 may communicate with a video streaming system 1100 via a local area network, virtual private network, or other network. In another alternative, the apparatus 1200 may be incorporated as a module or component of the system 1100 and communicate with other components via the bus 1214 or by some other modality.

[00138] The apparatus 1200 may be connected (e.g., via the bus 1214 and graphics processing unit 1220) to a display unit 1228. A display 1228 may include any suitable configuration for displaying information to an operator of the apparatus 1200. For example, a display 1228 may include or utilize a liquid crystal display (LCD), touchscreen LCD (e.g., capacitive display), light emitting diode (LED) display, projector, or other display device to present information to a user of the apparatus 1200 in a visual display.

[00139] One or more input devices 1230 (e.g., an alphanumeric keyboard, microphone, keypad, remote controller, game controller, camera or camera array) may be connected to the bus 1214 via a user input port 1222 to communicate information and commands to the

apparatus 1200. In selected embodiments, an input device 1230 may provide or support control over the positioning of a cursor. Such a cursor control device, also called a pointing device, may be configured as a mouse, a trackball, a track pad, touch screen, cursor direction keys or other device for receiving or tracking physical movement and translating the movement into electrical signals indicating cursor movement. The cursor control device may be incorporated into the display unit 1228, for example using a touch sensitive screen. A cursor control device may communicate direction information and command selections to the processor 1202 and control cursor movement on the display 1228. A cursor control device may have two or more degrees of freedom, for example allowing the device to specify cursor positions in a plane or three-dimensional space.

[00140] Particular embodiments may be implemented in a non-transitory computer-readable storage medium for use by or in connection with the instruction execution system, apparatus, system, or machine. The computer-readable storage medium contains instructions for controlling a computer system to perform a method described by particular embodiments. The computer system may include one or more computing devices. The instructions, when executed by one or more computer processors, may be configured to perform that which is described in particular embodiments.

[00141] As used in the description herein and throughout the claims that follow, “a”, “an”, and “the” includes plural references unless the context clearly dictates otherwise. Also, as used in the description herein and throughout the claims that follow, the meaning of “in” includes “in” and “on” unless the context clearly dictates otherwise.

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

CLAIMS

What is claimed is:

1 1. A method comprising:
2 receiving, by a computing device, user behavior for a user, wherein the user
3 behavior includes actions taken by the user on a video delivery service;
4 receiving, by the computing device, a real-time context based on the user
5 using the video delivery service;
6 generating, by the computing device, a set of action/entity pairs from a set of
7 actions for a set of entities found on the video delivery service based on the user behavior,
8 wherein each action/entity pair includes an action to be performed on an entity in which an
9 action can be taken on the video delivery system;
0 generating, by the computing device, probabilities for the set of action/entity
1 pairs based on the real-time context, wherein a probability for an action/entity pair indicates
2 the probability the user would select that action/entity pair when compared against other
3 action/entity pairs;
4 ranking, by the computing device, action/entity pairs in the set of action/entity
5 pairs based on respective probabilities for the action/entity pairs;
6 selecting, by the computing device, a list based on the ranking of the set of
7 action/entity pairs, the list including at least a portion of the set of action/entity pairs; and
8 outputting, by the computing device, the list to a client, wherein an action on
9 an entity in the list is performed when selected by the user.

1 2. The method of claim 1, further comprising:
2 training a first predictor using the user behavior, wherein the first predictor is
3 trained to output the set of action/entity pairs based on a probability the user would re-engage
4 with entities in the set of action/entity pairs.

1 3. The method of claim 1 or 2, wherein generating the set of action/entity
2 pairs comprises:
3 determining a first candidate set of entities the user has engaged within a first
4 time period using the user behavior;
5 determining if a probability the user would re-engage with each of the first
6 candidate set of entities is above a first threshold; and

7 including entities in the first candidate set of entities that have the probability
8 that meets the threshold in the set of action/entity pairs.

1 4. The method of claim 3, further comprising:
2 determining a second candidate set of entities the user has not engaged within
3 a second time period using the user behavior, the second candidate set of entities determined
4 to be related to entities the user engaged;
5 determining if the probability the user would engage with the second
6 candidate set of entities is above a second threshold; and
7 including entities in the second candidate set of entities that have the
8 probability that meets the threshold in the set of action/entity pairs.

1 5. The method according to any one of claims 1 to 4, further comprising:
2 training a second predictor using the user behavior, wherein the second
3 predictor is trained to rank action/entity pairs based on a probability the user would select an
4 action for an entity.

1 6. The method of claim 5, wherein training the second predictor
2 comprises:
3 training the second predictor using actions the user selected and actions for
4 entities the user did not select in the user behavior.

1 7. The method according to any one of claims 1 to 6, wherein the first
2 predictor generates an entity relationship graph that defines possible entities and relationships
3 between the entities.

1 8. The method of claim 7, wherein generating the set of action/entity
2 pairs comprises:
3 determining which of the entities in the entity relationship graph are available
4 for selection by the user on the video delivery service, and
5 selecting the set of action/entity pairs from the entities in the entity
6 relationship graph that are available for selection by the user.

1 9. The method according to any one of claims 1 to 8, wherein generating
2 probabilities for the set of action/entity pairs and ranking action/entity pairs in the set of
3 action/entity pairs comprises:

receiving first actions for manually added entities and second actions for engaged entities, the manually added entities being explicitly engaged by the user in the user behavior and the engaged entities being implicitly engaged by the user in the user behavior; and

ranking the first actions for manually added entities and the second actions for engaged entities into a set of third actions.

10. The method of claim 9, wherein generating probabilities for the set of action/entity pairs and ranking action/entity pairs in the set of action/entity pairs comprises:

receiving fourth actions for discovered entities other than the engaged entities and manually added entities, the discovered entities being determined based a relevance between the engaged entities and the different entities; and

ranking the fourth actions for the discovered entities.

11. The method of claim 10, wherein generating probabilities for the set of action/entity pairs and ranking action/entity pairs in the set of action/entity pairs comprises:

receiving fifth actions for a set of campaigns, the set of campaigns promoting campaign entities; and

ranking the fifth actions for the campaign entities in the set of campaigns.

12. The method of claim 11, wherein generating probabilities for the set of action/entity pairs and ranking action/entity pairs in the set of action/entity pairs comprises:

ranking the fifth actions based on a campaign urgency; and

ranking the fifth actions based on entity availability on the video delivery service for each campaign and the campaign urgency.

13. The method of claim 12, wherein generating probabilities for the set of action/entity pairs and ranking action/entity pairs in the set of action/entity pairs comprises:

receiving the ranked set of third actions and the ranked fourth actions; and

ranking the ranked set of third actions and the ranked fourth actions together into a set of sixth actions.

14. The method of claim 13, wherein generating probabilities for the set of action/entity pairs and ranking action/entity pairs in the set of action/entity pairs comprises:

receiving the set of sixth actions and the ranked fifth actions; and

4 ranking the set of sixth actions and the ranked fifth actions together into a set
5 of seventh actions, wherein the set of seventh actions is used to select the list.

1 15. The method according to any one of claims 1 to 14, wherein generating
2 probabilities for the set of action/entity pairs comprises using a function $f(u, e, a, c)$ to
3 generate the probabilities, where u includes the user behavior for the user, e is an entity, a is
4 an action, and c is the real-time context associated with the user.

1 16. The method according to any one of claim 1 to 15, wherein selecting
2 the list comprises:

3 filtering the set of action/entity pairs based on filter criteria to select the at
4 least the portion of the set of action/entity pairs.

1 17. The method according to any one of claim 1 to 16, further comprising:
2 receiving a selection of an action for an action/entity pair in the list; and
3 using a link for the action to perform the action on the video delivery service.

1 18. A non-transitory computer-readable storage medium containing
2 instructions, that when executed, control a computer system to be configured for:

3 receiving user behavior for a user, wherein the user behavior includes actions
4 taken by the user on a video delivery service;

5 receiving a real-time context based on the user using the video delivery
6 service;

7 generating a set of action/entity pairs from a set of actions for a set of entities
8 found on the video delivery service based on the user behavior, wherein each action/entity
9 pair includes an action to be performed on an entity found on the video delivery system;

10 generating probabilities for the set of action/entity pairs based on real-time
11 context, wherein a probability for an action/entity pair indicates the probability the user
12 would select that action/entity pair when compared against other action/entity pairs;

13 ranking action/entity pairs in the set of action/entity pairs based on respective
14 probabilities for the action/entity pairs;

15 selecting a list based on the ranking of the set of action/entity pairs, the list
16 including at least a portion of the set of action/entity pairs; and

17 outputting the list to a client, wherein an action on an entity in the list is
18 performed when selected by the user.




FIG. 1

2/14




FIG. 2

3/14




FIG. 3




FIG. 4




FIG. 5A

6/14

550 

552

554

556

Entity	Action	Probability
Show #1	1. Watch Latest Episode 2. Watch Next Episode 3. Watch Trailer 4. Watch Most Popular Episode	1. 70% 2. 90% 3. 30% 4. 50%
Show #2	Follow	60%
Movie #1	Watch	75%
Actor #1	Watch latest Episode	50%
Drama Genre	1. Watch New Release 2. Search For 3. Browse/select	1. 60% 2. 30% 3. 20%

FIG. 5B




FIG. 6A




FIG. 6B




FIG. 7

10/14




FIG. 8

11/14




FIG. 9

12/14




FIG. 10

13/14




FIG. 11

14/14




FIG. 12