



US 20220121697A1

(19) **United States**(12) **Patent Application Publication****Aravamudan et al.**(10) **Pub. No.: US 2022/0121697 A1**(43) **Pub. Date: Apr. 21, 2022**

(54) **METHODS AND SYSTEMS FOR SELECTING AND PRESENTING CONTENT BASED ON DYNAMICALLY IDENTIFYING MICROGENRES ASSOCIATED WITH THE CONTENT**

(71) Applicant: **Veveo, Inc.**, Andover, MA (US)

(72) Inventors: **Murali Aravamudan**, Andover, MA (US); **Ajit Rajasekharan**, West Windsor, NJ (US); **Kajamalai G. Ramakrishnan**, Nashua, NH (US)

(21) Appl. No.: **17/565,128**

(22) Filed: **Dec. 29, 2021**

Related U.S. Application Data

- (63) Continuation of application No. 17/231,394, filed on Apr. 15, 2021, which is a continuation of application No. 16/022,560, filed on Jun. 28, 2018, now Pat. No. 11,010,418, which is a continuation of application No. 15/360,592, filed on Nov. 23, 2016, now Pat. No. 10,037,326, which is a continuation of application No. 14/987,705, filed on Jan. 4, 2016, now Pat. No. 9,507,832, which is a continuation of application No. 14/803,412, filed on Jul. 20, 2015, now abandoned, which is a continuation of application No. 13/887,514, filed on May 6, 2013, now Pat. No. 9,092,503, which is a continuation of application No. 13/442,436, filed on Apr. 9, 2012, now Pat. No. 8,438,160, which is a continuation of application No. 12/844,366, filed on Jul. 27, 2010, now Pat. No. 8,156,113, which is a continuation of application No. 11/682,700, filed on Mar. 6, 2007, now Pat. No. 7,774,341.
- (60) Provisional application No. 60/834,966, filed on Aug. 2, 2006, provisional application No. 60/796,614, filed on May 1, 2006, provisional application No. 60/784,027, filed on Mar. 20, 2006, provisional application No. 60/779,547, filed on Mar. 6, 2006.

Publication Classification


- (51) **Int. Cl.**
G06F 16/438 (2006.01)
G06F 16/248 (2006.01)


G06F 16/28 (2006.01)
G06F 16/335 (2006.01)
G06F 16/435 (2006.01)
G06F 16/68 (2006.01)
G06F 16/951 (2006.01)
G06F 16/2453 (2006.01)
G06F 16/2457 (2006.01)
G06F 16/2458 (2006.01)
G06F 16/9535 (2006.01)
G06N 20/00 (2006.01)
H04N 21/442 (2006.01)
H04N 21/482 (2006.01)

- (52) **U.S. Cl.**
 CPC **G06F 16/438** (2019.01); **Y10S 707/99937** (2013.01); **G06F 16/284** (2019.01); **G06F 16/335** (2019.01); **G06F 16/337** (2019.01); **G06F 16/435** (2019.01); **G06F 16/437** (2019.01); **G06F 16/686** (2019.01); **G06F 16/951** (2019.01); **G06F 16/2453** (2019.01); **G06F 16/2457** (2019.01); **G06F 16/2462** (2019.01); **G06F 16/9535** (2019.01); **G06F 16/24578** (2019.01); **G06N 20/00** (2019.01); **H04N 21/44222** (2013.01); **H04N 21/4828** (2013.01); **Y10S 707/99943** (2013.01); **Y10S 707/99933** (2013.01); **Y10S 707/99934** (2013.01); **Y10S 707/99945** (2013.01); **G06F 16/248** (2019.01)

(57) ABSTRACT

A method of selecting and presenting content based on learned user preferences is provided. The method includes providing a content system including a set of content items organized by genre characterizing the content items, and wherein the set of content items contains microgenre metadata further characterizing the content items. The method also includes receiving search input from the user for identifying desired content items and, in response, presenting a subset of content items to the user. The method further includes receiving content item selection actions from the user and analyzing the microgenre metadata within the selected content items to learn the preferred microgenres of the user. The method includes, in response to receiving subsequent user search input, selecting and presenting content items in an order that portrays as relatively more relevant those content items containing microgenre metadata that more closely match the learned microgenre preferences of the user.








Figure 2




Figure 3




Figure 4




Figure 5





Figure 6






Figure 8




Figure 9




Figure 10




Figure 11




Figure 12




Figure 13




Figure 14




Figure 15

METHODS AND SYSTEMS FOR SELECTING AND PRESENTING CONTENT BASED ON DYNAMICALLY IDENTIFYING MICROGENRES ASSOCIATED WITH THE CONTENT

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation of prior U.S. patent application Ser. No. 14/987,705, filed Jan. 4, 2016, now U.S. Pat. No. 9,507,832, which is a continuation of U.S. patent application Ser. No. 14/803,412, filed Jul. 20, 2015, now abandoned, which is a continuation of U.S. patent application Ser. No. 13/887,514, filed May 6, 2013, now U.S. Pat. No. 9,092,503, which is continuation of U.S. patent application Ser. No. 13/442,436, entitled, Methods and Systems for Selecting and Presenting Content Based on Dynamically Identifying Microgenres Associated with the Content, filed Apr. 9, 2012, now U.S. Pat. No. 8,438,160, which is a continuation of prior U.S. patent application Ser. No. 12/844,366, entitled, Methods and Systems for Selecting and Presenting Content Based on Dynamically Identifying Microgenres Associated with the Content, filed Jul. 27, 2010, now U.S. Pat. No. 8,156,113, which is a continuation of U.S. patent application Ser. No. 11/682,700, entitled Methods and Systems For Selecting and Presenting Content Based On Dynamically Identifying Microgenres Associated With The Content, filed Mar. 6, 2007, now U.S. Pat. No. 7,774,341, which claims the benefit under 35 U.S.C. § 119(e) of the following applications:

[0002] U.S. Provisional Application No. 60/779,547, entitled A Framework for Learning User Behavior With Stochastic Signatures, filed Mar. 6, 2006;

[0003] U.S. Provisional Application No. 60/784,027, entitled A System And Method For Service Solicitation Enhanced With Relevant Personal Context to Elicit Targeted Response, filed Mar. 20, 2006;

[0004] U.S. Provisional Application No. 60/796,614, entitled A Learning Model For Multiple Dataspace With Applications To Mobile Environment, filed May 1, 2006; and

[0005] U.S. Provisional Application No. 60/834,966, entitled A Seminormalization Of Signatures For Reducing Truncation Errors And Stabilizing Relevance Promotion, filed Aug. 2, 2006;

the contents of all of which are hereby incorporated by reference herein in their entireties.

[0006] This application is related to the following applications, filed on Mar. 6, 2007:

[0007] U.S. patent application Ser. No. 11/682,693, entitled Methods and Systems For Selecting and Presenting Content Based On Learned Periodicity Of User Content Selection, now U.S. Pat. No. 7,774,294;

[0008] U.S. patent application Ser. No. 11/682,689, entitled Methods and Systems For Selecting and Presenting Content Based On Activity Level Spikes Associated With The Content, now U.S. Pat. No. 7,657,526;

[0009] U.S. patent application Ser. No. 11/682,695, entitled Methods and Systems For Selecting and Presenting Content Based On User Preference Information Extracted From An Aggregate Preference Signature, now U.S. Pat. No. 7,739,280;

[0010] U.S. patent application Ser. No. 11/682,533, entitled Methods and Systems For Selecting and Pre-

senting Content Based On A Comparison Of Preference Signatures From Multiple Users;

[0011] U.S. patent application Ser. No. 11/682,596, entitled Methods and Systems For Segmenting Relative User Preferences Into Fine-Grain and Coarse-Grain Collections, now U.S. Pat. No. 7,529,741;

[0012] U.S. patent application Ser. No. 11/682,588, entitled Methods and Systems For Selecting and Presenting Content On A First System Based On User Preferences Learned On A Second System, now U.S. Pat. No. 7,835,998; and

[0013] U.S. patent application Ser. No. 11/682,599, entitled Methods and Systems For Selecting and Presenting Content Based On Context Sensitive User Preferences, now U.S. Pat. No. 7,792,815.

BACKGROUND

Field of Invention

[0014] This invention generally relates to learning user preferences and, more specifically, to using those preferences to personalize the user's interaction with various service providers and interactions with content query systems, e.g., to better find results to queries provided by the user and to ordering the results for presentation to the user.

Description of Related Art

[0015] Personalization strategies to improve user experience can be chronologically classified into two categories: (1) collaborative filtering and (2) content reordering. Each is summarized in turn.

[0016] Collaborative Filtering was used in the late 1990s to generate recommendations for users. The term collaborative filtering refers to clustering users with similar interests and offering recommendations to users in the cluster based on the habits of the other users. Two distinct filtering techniques—user based and item based—are used in filtering.

[0017] In U.S. Patent App. Pub. No. U.S. 2005/0240580, Zamir et al. describe a personalization approach for reordering search queries based on the user's preferences. The application describes a technique for learning the user's preferences and increasing the promotion level of a search result based on personalization. Zamir et al. create a user profile, which is a list of keywords and categories listing the user preferences. The profile is generated from multiple sources, such as (1) information provided by the user at the time the user registers a login, (2) information from queries that the user has submitted in the past, and (3) information from web pages the user has selected.

[0018] Some systems directed to reordering content in the context of television schedules define categories and sub-categories according to an accepted standard. User preferences are gathered using various models such as (1) user input, (2) stereotypical user models, and (3) unobtrusive observation of user viewing habits. In some implementations, these models operate in parallel and collect the user preference information.

[0019] In other systems, a set of fixed attributes is defined and all media content and all user preferences are classified using these attributes. A vector of attribute weights captures the media content and the user preferences. The systems then determine the vector product between the content

vector and the user preferences vector. The system suggests content to users where the values of the vector products exceed a predetermined threshold.

BRIEF SUMMARY

[0020] The invention provided methods and systems for selecting and presenting content based on learned user preferences.

[0021] Under an aspect of the invention, a user-interface method of selecting and presenting a collection of content items in which the presentation is ordered at least in part based on learning the preferred microgenres of content of the user as contained in content items selected by the user includes providing a content system including a set of content items organized by genre information that characterizes the content items, wherein the genre information is specified by the content system, and wherein the set of content items contains microgenre metadata further characterizing the content items. The method also includes receiving incremental input entered by the user for incrementally identifying desired content items and, in response to the incremental input entered by the user, presenting a subset of content items to the user. The method further includes receiving actions from the user selecting content items from the subset and analyzing the microgenre metadata within the selected content items to learn the preferred microgenres of the user. The method includes, in response to receiving subsequent incremental input entered by the user, selecting and presenting a collection of content items in an order that portrays as relatively more relevant those content items containing microgenre metadata that more closely match the learned microgenre preferences of the user.

[0022] Under another aspect of the invention, at least one learned microgenre preference is stored on a user client device and selecting and presenting the collection of content items includes selecting and presenting content items stored on the client device.

[0023] Under a further aspect of the invention, the method also includes analyzing the genre information characterizing the selected content items to learn the preferred genres of the user, and wherein the selecting and presenting the collection of content items is further based on promoting the relevance of those content items characterized by genre information that more closely match the learned genre preferences of the user.

[0024] Under yet a further aspect of the invention, the method further includes analyzing the date, day, and time of the user selection actions and analyzing at least one of the genre information and microgenre metadata of the selected content items to learn a periodicity of user selections of similar content items. The similarity is determined by comparing the at least one of the genre information and microgenre metadata of the selected content item with the previously selected content item. The periodicity indicates the amount of time between user selections of similar content items relative to a reference point. The method also includes associating the learned periodicity with the at least one of the genre information and microgenre metadata of the similar content items. The selecting and presenting the collection of content items is further based on promoting the relevance of those content items characterized by genre information or containing microgenre metadata associated with periodicities similar to the date, day, and time of the subsequent incremental input.

[0025] Under yet another aspect of the invention, the selecting and ordering the collection of content items is further based on popularity values associated with the content items. Each popularity value indicates a relative measure of a likelihood that the corresponding content item is desired by the user.

[0026] Under a further aspect of the invention, the set of content items includes at least one of television program items, movie items, audio/video media items, music items, contact information items, personal schedule items, web content items, and purchasable product items. The descriptive terms can include at least one of title, cast, director, content description, and keywords associated with the content.

[0027] Under yet a further aspect of the invention, the set of content items is contained on at least one of a cable television system, a video-on-demand system, an IPTV system, and a personal video recorder.

[0028] These and other features will become readily apparent from the following detailed description wherein embodiments of the invention are shown and described by way of illustration.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0029] For a more complete understanding of various embodiments of the present invention, reference is now made to the following descriptions taken in connection with the accompanying drawings in which:

[0030] FIG. 1 illustrates two modules of a leaning engine.

[0031] FIG. 2 illustrates a collections of signatures of a user.

[0032] FIG. 3 illustrates a collection of signatures of a user for a single dataspace in a single location.

[0033] FIG. 4 illustrates orthogonal periodicities.

[0034] FIG. 5 illustrates overlapping periodicities.

[0035] FIG. 6 illustrates sample vectors in three vector spaces.

[0036] FIG. 7 illustrates seminormalization of signature probabilities.

[0037] FIG. 8 illustrates an example of a seminormalized signature.

[0038] FIG. 9 illustrates an example of detecting an increased level of activity associated with a content item.

[0039] FIG. 10 illustrates a context-specific personal preference information service.

[0040] FIG. 11 illustrates the local tracking and strengthening of the personal preference signatures based on user activity and the content on a mobile device.

[0041] FIG. 12 illustrates the information flow when a user device makes a request to a service provider.

[0042] FIG. 13 illustrates an alternate information flow when a user device makes a request to a service provider.

[0043] FIG. 14 illustrates examples of services that benefit from the context-sensitive personal preference service.

[0044] FIG. 15 illustrates possible user device configurations for use with the learning system and the context-sensitive personal preference service.

DETAILED DESCRIPTION

[0045] Preferred embodiments of the invention capture user preferences for a single user or a family of users based on historical observations of the users' activities and by use

of a statistical learning model (also referred to as a learning engine). In an illustrative embodiment, the users of the family are members of a household using a single interface device. The learning model identifies a signature, or set of signatures, for the entire household as a stochastic signature. This stochastic signature is used to predict future activities in which members of the family may engage. For the sake of simplicity, the description will refer to signatures and activities of a single user, when in fact, the same applies to a group of users using a single interface device.

[0046] One benefit of the learning engine is to enhance the user's experience by increasing the accuracy of the results of a user's query for content and organizing these results so as to put the most likely desired content items at the top of the results list. This increases the user's speed and efficiency in accessing the content he or she desires. In addition, the signatures can be used to identify clusters of users with similar interests for use with collaborative filtering techniques. The learning system can be one aspect of an overall system that provides content and selectable actions to the user, or the learning system can be a standalone system that monitors the user's actions across multiple systems and provides learned user preferences to external systems based on the context of the user's interaction with a particular external system.

[0047] Information Captured by Signatures

[0048] The stochastic signature is a probabilistic model, and it is an identity that can be used for validation, prediction, etc. While this type of signature can be used to identify a user according to his or her preference, it is distinguished from a unique signature that a user may employ to authenticate an item or guarantee the source of the item, for example. A stochastic signature may be created based on various types of activities, for example, watching television programs, dialing telephone numbers, listening to music, etc. Thus, embodiments of the invention are useful in a wide variety of contexts. In applications where there are multiple dataspace to be searched, the system will use a collection of stochastic signatures related to that dataspace. For example, when searching a personal address book, calendar, or air-line schedules, the system can use the set of signatures related to that particular dataspace. In addition, the system can also learn correlated activities across dataspace. This allows the system to learn how the user interacts with related dataspace and use that learned behavior to personalize the presentation of content to the user. However, for the sake of simplicity, certain embodiments of the invention will be described with reference to a single dataspace interface, e.g., a television system interface.

[0049] In the context of a user device with limited input capability, for example, a television remote control, the stochastic signature is particularly useful because it can be difficult and/or time consuming to enter queries on such a limited input device. The stochastic signature increases the likelihood that the desired search results will be found based on limited query input. For example, if a particular household has generally watched a certain program at a given time in the past, stochastic signatures can be used to predict that the household will watch the program at the given time in the future. Thus, instead of requiring a member of the household to enter the title of the program, the learning system can predict that the member wishes to watch the program based on only a few button presses.

[0050] Embodiments of the present invention build on techniques, systems and methods disclosed in earlier filed applications, including but not limited to U.S. patent application Ser. No. 11/136,261, entitled Method and System For Performing Searches For Television Programming Using Reduced Text Input, filed on May 24, 2005, U.S. patent application Ser. No. 11/246,432, entitled Method And System For Incremental Search With Reduced Text Entry Where The Relevance Of Results Is A Dynamically Computed Function of User Input Search String Character Count, filed on Oct. 7, 2005, and U.S. patent application Ser. No. 11/235,928, entitled Method and System For Processing Ambiguous, Multiterm Search Queries, filed on Sep. 27, 2005, the contents of which are hereby incorporated by reference. Those applications taught specific ways to perform incremental searches using ambiguous text input and method of ordering the search results. The present techniques, however, are not limited to systems and methods disclosed in the incorporated patent applications. Thus, while reference to such systems and applications may be helpful, it is not believed necessary to understand the present embodiments or inventions.

[0051] FIG. 1 shows the architecture of an illustrative learning engine 100. There are two distinct modules to learning engine 100—a data collection module 105 and a signature computation module 110. Data collection module 105 monitors the user activity for channel tuning, DVR recording, etc. and captures the relevant statistics of the activity, for example, the duration a TV channel was watched, as well as the genres and microgenres (discussed below) of the program that was watched. In the case of a mobile device, additional information is collected, such as the type of dataspace being visited (e.g., phone book, calendar, and downloadable media content), geographic location of the mobile device, etc. Data collection module 105 can reside in a client device, where it gathers data about the users' activities and sends this data to signature computation module 110. In the alternative, data collection module 105 can reside on a remote server that serves content to the client device. In this case, the remote server collects data about the content requested by the users and passes this data to computation module 110.

[0052] As mentioned, the learning engine gathers information about channels, genres, and microgenres that the user has watched. Herein, the term “channel” refers to a tunable entity in a television system. A channel can be identified by its name (CBS, ABC, CNBC, etc.).

[0053] The term “genre” refers to the overall theme of an item. In some systems, every retrievable item is categorized into a genre. The collection of genres is system-definable, and can be as coarse or as fine-grained as necessary. In addition, the genres can be defined independent of the retrievable items and can be defined ahead of populating a content system with retrievable items. In one implementation, a function $g(x)$ returns a subset of the set of genres for a given item. Thus, $g(x)$ is a function with a domain space of a set of retrievable items and the range space of a set of all subsets of genres. This is so because any retrievable item may belong to more than one genre, e.g., a movie *Sleepless in Seattle* has a genre of movie and romance.

[0054] The term “microgenre” refers to a very refined, unambiguous theme of descriptor for a given item. For example, New England Patriots as a search item has a microgenre of NFL Football and genres of football and

sports. As with genres, a search item can have multiple microgenres. While genres are “macro” themes, microgenres are “micro”, unambiguous themes; these themes come from descriptive terms and metadata within the search items. Thus, the microgenres for New England Patriots also include Tom Brady. Microgenres are not limited to a set of predetermined descriptors, as are genres in the prior art, but can be any word that describes the item. Whether a particular theme is a genre or microgenre depends on the particular item with which the theme is associated and the configuration of the content system. Thus, microgenres are dynamic and generated “on-the-fly”, while genres are static and system defined.

[0055] In dataspace other than the television content space, the channel, genre, and microgenre approach to characterizing items is modified to reflect the attributes of the content items in that particular dataspace. Thus, for a telephone directory dataspace, the channel statistics are replaced with statistics related to the person or entity called. The genre statistics are replaced by statistics related to the type of entity called, for example, individual or business. While the microgenre statistics are replaced by statistics related to key secondary attributes of the item, such as home, office, and mobile telephone numbers as well as, for example, telephone numbers of persons related to the persons called.

[0056] Computational module **110** is sent the current day’s data periodically and determines the users’ signatures. In so doing, computational module **110** combines this current data with historical data using exponential smoothing or other smoothing techniques (discussed below) so that the signatures adapt over time to the users’ changing preferences. Computational module **110** also performs other computations involving the signatures, for example, combining the individual signatures to obtain aggregate signatures that predict the viewing preferences of a large collection of individuals or creating signatures that capture the average activity level associated with a particular program (described in greater detail below). In one embodiment of the system, computational module **110** resides in one or more servers, to exploit the computational power of larger processors. However, in some implementations, e.g., where privacy is an issue, computational module **110** may reside in the client device.

[0057] A particular stochastic signature is a normalized vector of probabilities. The probabilities capture the historical patterns of the user’s behavior with respect to a particular set of activities. An example of a signature for use with a television system is {(ABC 0.3), (CBS 0.2), (NBC 0.5)}. This signature captures that over a given time period, when the user was watching television, the user watched ABC 30% of the time, CBS 20% of the time, and NBC 50% of the time. The stochastic nature of the signature says that this is a historical average and not an exact number.

[0058] Because the system captures the user’s behavior and preferences across multiple dataspace, taking into account the geographic location of the user, or the user’s client device, the multiple signatures can be represented as a set with three indices. Thus, the convention signature(t, g, s) represents a signature in geographic location g at time t for dataspace s. This allows the system to use difference subspace projections to utilize the information contained in the entire set. For example, the system may utilize the user’s preferences based on activity across all geographic locations

or based on a composite of multiple times for a given dataspace and given location. The composite signature is described in greater detail below.

[0059] Although time is obviously a continuous variable, for the purpose of learning the user’s preferences and activities, a coarse, or discretized, view of time is used to capture all activity. Thus, the system divides time into discrete quantities and represents time as an integer from one to the number of discrete quantities in a day. For example, time can be divided into the number of minutes in a day, thereby representing time as a number 1 to 1440. In addition, this discrete representation of time can be further subdivided into time slots that encompass multiple minutes, as discussed below. The duration, and therefore the number, of time slots is selected based on the nature of the activity of the particular dataspace. For example, in the television dataspace it is appropriate to divide time into 30-minute time slots to correspond to the program boundaries. In other dataspace, the coarseness can vary. Although it is not necessary to have the same time division in all dataspace, the examples set forth below assume identical time slot durations for the sake of simplicity. Similarly, geographic location, though continuous, is discretized and represented by character strings. For example, the geographic location identifiers can be a postal code, a major metropolitan area, or an area of a given size with the latitude and longitude of its center being the location identifier.

[0060] There are many possible collections of signatures that capture the activities of the user or family of users at various levels of granularity. FIG. 2 shows a sample signature hierarchy for the multiple dataspace learning model, with n locations **200**, m dataspace **210**, and k time slots **220**. At the first level, the figure illustrates activities in each location **200**. Within each location **200**, the system captures dataspace-specific activities in individual signatures. Inside each dataspace **210**, for each time slot **220**, the system obtains a unique signature. Finally, the signature hierarchy captures the nature of the activity within the time slot by appropriate keyword **230**, genre **240**, and microgenre signatures **250** (or equivalent statistics depending on the dataspace, as described above). The illustrative learning system shown in the figure has 3nmk signatures in the collection.

[0061] The timeslots shown in FIG. 2 can be further divided according to the particular needs of the learning system. Thus, a top-level time slot can have lower level time slots organized beneath the top-level time slot. For example, a top-level time slot can be a day organized into lower-level time slots of an hour or half-hour increments, each having its own collection of signatures. Similarly, the day time slot can have a collection of composite signatures beneath it that aggregate all of the information of the individual time slots for that given day into a single composite signature.

[0062] FIG. 3 shows an illustrative example of the organization of a signature collection **300** for the user in the television program dataspace at a single location. At the top level, the signatures are classified into various periodicities for the user, as discussed in greater detail below. The example in FIG. 3 shows a weekday periodicity **305** and a weekend periodicity **310**. Within each periodicity, signature collection **300** is further divided into individual time slots **315** with a composite **320** for each day. Within each further division exists three types of signatures—channel **325**, genre **330**, and microgenre **335**. Thus, there is one of each of these

three types of signatures for every weekday time slot, weekend time slot, and one for each weekday composite and weekend composite. Therefore, the system captures the activities performed by the user in this single dataspace and at this single location as defined by the hierarchy present in signature collection 300.

[0063] Because activities vary widely in a multiple dataspace environment, the system can capture the user's activities, according to the signature hierarchy, as duration and/or count. In other words, the system can track the amount of time the user spent performing an activity, or the system can track the number of times the user performed a particular activity, or the system can record both. For example, if the system is modeling a DVR recording activity or DVD ordering activity, there is no duration measure associated with it. Thus, in these cases, the system will capture the intensity of the activities by the count (frequencies) of the activities. However, other activities have duration as a natural measure of the intensity of the activities (e.g., watching a television program). While still other activities have both count and duration as a natural measure of intensity (e.g., placing a telephone call and conducting the call). To be inclusive of all activities, the system models every activity by both count and duration. Thus, there are two signatures for each keyword, genre, and microgenre division of the hierarchy. Likewise, there are two for each composite as well. For each time, location, and dataspace, a function/defines the convolution of the two intensity measures into a single signature:

$$figs: (\text{count}, \text{duration}) \rightarrow \text{single measure} \quad (\text{Equation 1})$$

[0064] For the sake of simplicity, this description omits the adjective, count or duration, in referring to signatures, opting for disambiguation based on the context.

[0065] In one embodiment of the invention, signatures capture the television viewing activity of the family of users in a single geographic location, and these signatures are used to identify and organize television program search results. The learning engine divides a day into time slots. For the purposes of this example, there are 48 time slots in a 24-hour day. Thus, one time slot corresponds to the smallest length TV-program, i.e., 30 minutes. In other implementations, time slots may be larger or smaller, for example, in a system using stochastic signatures to identify a user's telephone calling preferences, the time slots may be two to three hours long. During each time slot, the user activity is recorded and the learning system creates a time slot signature for each time slot. In addition, at the end of each day, the learning system creates a composite signature based on the data collected across all time slots within the current day. The signature is said to be composite in that it represents a user's activity across multiple time slots. As discussed in greater detail below, the learning system uses smoothing techniques to create evolving signatures that retain activities in the distant past as well as the most recent activities.

[0066] The day is divided into time slots because each family of users has a recurring viewing behavior based on the time of day. Thus, the learning system learns television-viewing preferences from the past behavior during a given time slot. Any queries for television content that originate in

that time slot on a later day can use these preferences to identify and organize content by using that time slot's signature.

[0067] For example, in an illustrative family of three—husband, wife, and a child—mornings and afternoons are taken up by soap operas and talk shows; afternoons are taken up by cartoons and children's programming; and evenings are taken up by business, movies, and prime time shows. During these periods, it is likely that queries for current television content also relate to the corresponding past viewing behavior. Thus, signatures that capture this past behavior are used to identify and organize content consistent with this past behavior. However, for more aggregate behavior, independent of time slots, it is desirable to have a coarse grain view of the day's activity in the household. The time slot activity is aggregated into a day's activity; this is the basis of the composite signature. Thus, at the end of each day, the system has collected and aggregated 49 different signatures, 48 individual time slot signatures and one composite signature.

[0068] Composite signatures serve two purposes. First, if the family of users has a slight time-drift in their behavior (e.g., some days a particular user watches a given program at 10:00 AM, while other days at 10:45 AM), the time slot signatures may get shifted by one slot. However, the composite will still capture the overall viewing behavior correctly. Second, a particular user may time-shift deliberately by many time slots. The composite signatures will also correctly capture this behavior.

[0069] User Periodicity

[0070] The above example implicitly assumes that the user has a recurring behavior with a periodicity of a day. However, the learning system may utilize other periodicities, as explained below. As mentioned above, one benefit of the learning system is to enhance the user's experience by increasing the accuracy of the results of a user's query for content and organizing these results so as to put the most likely desired content items at the top of the results list. This increases the user's speed and efficiency in accessing the content he or she desires.

[0071] Towards this end, the learning system infers periodicity of activities. For example, as discussed above, there is a daily periodicity of activities. However, the daily periodicity model may not always apply, as occurs during a weekend, during which time the users' activities can be quite different from those during the week. To capture this different behavior pattern, for example, the system will utilize two different periodicities. Thus the weekday periodicity will contain data for the days during the week, while the weekend periodicity will be empty for those days and vice versa. This is an example of orthogonal periodicity.

[0072] The term orthogonal refers to the fact that the periodicity waveforms are orthogonal; i.e., if $f(x)$ is a periodicity function and $g(x)$ is another periodicity function, then $f(x)$ and $g(x)$ have orthogonal periodicity if

$$f(x)g(x) = 0; \quad (\text{Equation 2})$$

$$0 \leq x \leq \infty$$

[0073] Equation 2 defines strong orthogonality, or point-wise-orthogonality, in contrast with the strict mathematical definition of orthogonality of functions (see F. B. Hildeb-

rand, *Introduction to Numerical Analysis*, second edition, McGraw-Hill Book Company, New York, 1974, hereby incorporated by reference). FIG. 4 illustrates an example of orthogonal periodicity. The figure shows a variety of waveforms that represent the activity captured by a set of signatures for a particular dataspace in a particular location. The Y-axis is the intensity of activity during a particular day, and X-axis is the day. A weekday waveform 405 captures the activity during the days of the week (i.e., Monday-Friday). Whereas a Saturday waveform 410 and a Sunday waveform 415 capture the activity on Saturday and Sunday, respectively. A solid line shows weekday periodicity waveform 405; a short dashed line shows the Saturday periodicity waveform 410; and a long dashed line show Sunday periodicity waveform 415.

[0074] FIG. 4 illustrates the waveforms are orthogonal in that the activity level for weekday waveforms is zero during Saturday and Sunday, while the Saturday waveform is zero for all non-Saturday days, and the Sunday waveform is zero for all non-Sunday days. The system captures these orthogonal periodicities by storing the activities in distinct sets of signatures, with one set of signatures for each orthogonal period. As explained above, the set can include both individual time slot signatures as well as a composite signature for the entire orthogonal period. When a user query is submitted within a particular period, the corresponding set of signatures is used in identifying and organizing the search results.

[0075] Although the above example is in terms of a week, periodicity can extend beyond a week. For example, periodicities can exist within a day or can extend beyond a week. In addition, some periodicities may not be orthogonal. Thus, the system uses a second kind of periodicity, namely overlapping periodicity, to capture this phenomenon.

[0076] In overlapping periodicities, the periods overlap; i.e., the same time and day can belong to multiple periods, one having a larger frequency than the other. Thus, the strong orthogonality property of Equation 2 does not apply to overlapping periodicities. FIG. 5 shows an example of overlapping periodicity. In this example, a user watches a recurring program every Wednesday, along with the usual daily programs that she watches. Thus, there is a weekly period 505 with a frequency of once per week and a daily period 510 with a frequency of once per day.

[0077] Overlapping periodicities are distinguished by storing the same activities in multiple sets of signatures, one set for each overlapping periodicity. In the example of FIG. 5, the system will store the same Wednesday activity both in daily set 510 and weekly set 505. Notice that weekly set 505 does not contain activities from other days. When a query is submitted on a Wednesday, a union of both signatures is used in identifying and organizing the content results. Both signatures are combined in such as way as to reflect the fact that the weekly signature 505, anchored on Wednesdays, has a greater impact on the query results than does daily signature 510.

[0078] As mentioned above, the learning system defines periodicities according to the users' behavior. In one illustrative implementation, the system compares each recorded user action and determines the periodicity of similar actions by measuring the time that elapses between the similar actions. The similarity can be based on, for example, the user watching the same television channel or the user watching television programs of the same genre or microgenre. There-

fore, if a user watches a particular television show on the first Tuesday of every month, the system would capture this as a monthly periodicity. Thus, although the system can use predefined periodicities, the system creates periodicities of any time duration as needed by the particular usage case. As mentioned above, capturing the user's viewing preferences in the television dataspace is only one example of an implementation of the learning system. Other examples include learning the user's dialing preferences in the telephone dataspace or tracking the user's buying behavior in the internet dataspace.

[0079] Signatures as Multiple Vectors

[0080] As explained above, vectors can be used to capture the family of users' behavior. A vector is defined as a linear array of numbers. Every vector belongs to a vector space. In one embodiment of the learning system, the system only operates in vector spaces in R^{+n} , defined as

$$R^{+n} = \{(x_1, x_2, \dots, x_n) \mid x_i \geq 0 \text{ all } i\} \quad (\text{Equation 3})$$

[0081] The dimensionality of the vector space depends on the type of vector. For example, the dimensionality of a channel vector space is the number of channels in the television system. The values in the vector also depend on the type of vector; e.g., it can be duration or count or any other metric deemed appropriate to capture the activity of the family of users. FIG. 6 shows an example of a collection of vectors capturing the users' activity between 10:00 AM and 10:30 AM on a weekday.

[0082] The vectors in FIG. 6 correspond to three different vector spaces—channel, genre, and microgenre. The dimensions of these vector spaces are the number of channels in the TV system, the number of genres defined in the learning system, and the number of microgenres dynamically created in the learning system, respectively. Only nonzero values are stored in the various vectors. All other values are implicitly zero and are not stored in the system. Thus, the learning system fundamentally stores all vectors as sparse vectors. The technology of sparse vectors and sparse matrix computations eases the burden of working with large vector spaces (see I. S. Duff, A. M. Erisman, and J. K. Reid, *Direct Methods for Sparse Matrices*, Monographs on Numerical Analysis, Oxford Science Publications, Clarendon Press, Oxford, 1986, for a description of numerical computations using sparse matrices and sparse vectors, hereby incorporated by reference).

[0083] A channel vector 605 in the figure has nonzero values for channels CBS, ABC, and CNBC. The values correspond to the number of minutes the user watched each channel between 10:00 AM and 10:30 AM. Similarly, the program genres are captured in a genre vector 610. In this example, the CBS and CNBC channels were broadcasting programs of type business and ABC was broadcasting a program of type comedy. Finally the program microgenres are captured in a microgenre vector 615. In the above example, ABC was broadcasting the comedy show *Seinfeld*, CNBC was broadcasting a business show *Squawkbox*, and no microgenre was created for the CBS show.

[0084] As previously mentioned, the techniques described above can be implemented in data collection modules and signature computation modules that reside on either a client device or a remote server system. Thus, the channel, genre,

and microgenre data can be gathered and processed locally by the client device, or this information can be sent to a remote server system for processing. Likewise, the signatures can reside on a client device or on a remote server system for use as described below.

[0085] In addition to capturing the user's activities according to keyword (i.e., channel in the television dataspace context), genre, and microgenre, the system also learns the amount of time the user spends in each dataspace independent of location and time slot. This gives rise to yet another signature: the dataspace fraction signature. The dataspace fraction signature (herein "dfs") has the coordinates of time and location and is represented by $dfs(t, g)$. The signature $dfs(t, g)$ is a normalized probability vector indicating the fraction of time (and/or activity count) the user spent in various dataspace. For example, $dfs(t, g)[s]$ contains the value indicating the fraction of time and/or count spent in dataspace s , at time t in location g . This two-coordinate signature is used to reorder the search results space when a search across dataspace is performed. Meanwhile, as described above, a three-coordinate signature is used to reorder the items within each dataspace, e.g., $ks(t, g, s)$ denotes a keyword signature in time slot t , in location g , and in dataspace s ; $ks(t, g, s)[x]$ denotes the value of element x in the keyword signature. Therefore, when the user initiates a search across all dataspace, the system will reorder content items from the multiple dataspace according to the user's dataspace preferences based on the information contained in the dataspace fraction signature. If the user performs actions in one particular dataspace relative to another, the result from the more heavily used dataspace would be promoted over the results from the lesser-used dataspace.

[0086] The following example is provided to illustrate this aspect of the learning system. A mobile user visited the telephone dataspace 30 times, the television dataspace 20 times, and the web dataspace 10 times while located in Denver during the LOAM time slot. During these interactions with the system, the user called Sam, Sally, and Stewart, speaking for 5, 15 and 10 minutes respectively. The user watched a television program entitled "Seinfeld" for 30 minutes. In addition, the user browsed the Google webpage for 10 minutes and Yahoo! webpage for 20 minutes, respectively. Using a count measure for the dataspace fraction signature and a duration measures for the television, telephone, and web dataspace, the keyword signature and dataspace fraction signature ensemble, will be as follows:

$$\begin{aligned}
 dfs(10, \text{"Denver"})[\text{"phone-space"}] &= 0.5 \\
 dfs(10, \text{"Denver"})[\text{"TV-space"}] &= 0.33 \\
 dfs(10, \text{"Denver"})[\text{"web-space"}] &= 0.17 \\
 ks(10, \text{"Denver"}, \text{"phone-space"})[\text{"Sam"}] &= 0.17 \\
 ks(10, \text{"Denver"}, \text{"phone-space"})[\text{"Sally"}] &= 0.50 \\
 ks(10, \text{"Denver"}, \text{"phone-space"})[\text{"Stewart"}] &= 0.33 \\
 ks(10, \text{"Denver"}, \text{"TV-space"})[\text{"Seinfeld"}] &= 1.0 \\
 ks(10, \text{"Denver"}, \text{"web-space"})[\text{"Google"}] &= 0.33 \\
 ks(10, \text{"Denver"}, \text{"web-space"})[\text{"Yahoo!"}] &= 0.67
 \end{aligned}$$

[0087] Thus, if the user enters a text query starting with the letter "S", all results beginning with the letter "S" would be presented to the user. However, the matching results from the phone-space would be promoted over the results from the TV-space and the web-space because the dataspace fraction signature probability for the phone-space is the greatest. This is so even though the probability for the lone TV-space item is greater than any of the phone-space items. Within the phone-space, the individual items would be sorted according to the keyword signature probability values. Therefore, the entry for "Sally" would be promoted over the other phone-space items. This example clearly shows the Bayesian property of the signatures. That is, the probabilities add up to one, conditioned on the fact that the user is in a particular dataspace (see B. W. Lindgren, G. W. McElrath, D. A. Berry, *Introduction to Probability and Statistics*, Macmillan publishing co., New York, 1978, herein incorporated by reference, for more details on Bayes's theorem).

[0088] As described above, the signatures associated with a particular dataspace (i.e., keyword, genre, and microgenre signatures) capture the probability of the user performing a future action or desiring a future content item based on past activities and selections that took place within that particular dataspace. Thus, the individual dataspace signatures are conditional signatures in that the probabilities they measure are conditioned upon the user operating in the particular dataspace to which those signatures relate.

[0089] The dataspace fraction signature probabilities can be used to weight the individual dataspace signature probabilities to provide a measure of the probability of a user performing a given action outside of a particular dataspace. This operation gives rise to an unconditional signature. The unconditional signature measures the probability of the user performing an action outside of any particular dataspace based on the information contained in the individual dataspace signatures and the information contained in the dataspace fraction signatures. The system uses Equation 4, below, to determine the unconditional keyword signature for an activity "A". Unconditional signatures for genre and microgenre can be determined in the same way.

$$uks(t, g, s) = ks(t, g, s)[A] * dfs(t, g)[s] \quad (\text{Equation 4})$$

[0090] The learning system can organize the various dataspace, content, and selectable actions into a tree hierarchy, which the user can navigate using the client device. The unconditional signatures are used by the learning system to rearrange the various branches of the tree structure so as to present the most favored content to the user based on the probabilities stored in the signatures. In addition, the unconditional probabilities enable the system to present lists of commingled selectable actions and content items based on the most commonly performed actions and most commonly exhibited preferences. For example, the learning system is capable of creating a "My Favorites" list based on the various signatures, or the system could rearrange a content tree hierarchy in order to reduce the effort required of the user to reach certain preferred content.

[0091] Correlated Activity Signatures

[0092] The learning system is also capable of learning correlated activities across dataspace. A correlated activity is an activity performed in a secondary dataspace while

starting from a primary dataspace. In general, by capturing correlated activities across dataspace, the system is learning not only standalone actions and content preferences, but the system is learning chains of actions performed by the user. For example, a user enters the telephone dataspace of his device to make a telephone call. During the telephone call, the user wishes to enter the calendar dataspace to search for and review the date and time of a particular appointment. In this example, the user remains engaged with the primary dataspace, the telephone dataspace, for the duration of the telephone call. The user also performs a correlated action, the act of searching for an appointment, in the secondary dataspace, which is the calendar dataspace.

[0093] The purpose of learning the correlated activities is to achieve better ordering of the search query results in the secondary dataspace based on the correlated activities learned by the system. Thus, the correlated activity signatures provide yet another way to learn the preferences of the user and how the user interacts with his or her client device. This additional set of preferences and learned actions further enhances the user experience.

[0094] In general, the system has an activity matrix A that is a square N by N matrix, where N is the number of dataspace. Each entry in the matrix is a signature vector that captures the actions performed in the secondary dataspace while operating from the primary dataspace. Thus, A is in fact a three dimensional matrix, which can be defined as follows:

$$A(i, i)[x] := 0; 1 \leq i \leq N, \text{ for all items } x \in \text{dataspace } i \quad (\text{Equation 5})$$

$$A(i, j)[x] :=$$

$$\text{average number of accesses of item } x \text{ in dataspace } j$$

$$\text{while in dataspace } i; 1 \leq i \leq N; 1 \leq j \leq N; i \neq j;$$

$$\text{for all items } x \in \text{dataspace } j$$

[0095] The matrix determined by Equation 5 captures the correlated activities of the user, and therefore can be used in accordance with the techniques disclosed herein to predict the probability that the user would perform an action in a secondary dataspace while operating from a primary dataspace. In addition, the correlated activity signatures can be used to determine the unconditional probability of the user accessing a keyword item x in dataspace s , location g , and at time t . The probability determination depends in part on the mode of access utilized by the user. As described in Equation 6 below, if the user is entering dataspace s at the root level of the client device (i.e., the user is not currently in a dataspace), the probability determination is based on the dataspace fraction signature and the relevant signatures for the selected dataspace (e.g., the keyword signature, the genre signature, or the microgenre signature). If the user is entering dataspace s from another dataspace, the probability determination is based on the dataspace fraction signature and the correlated activity matrix A .

$$Prob[x] = \quad (\text{Equation 6})$$

$$\left\{ \begin{array}{l} dfs(t, g)[s] * ks(t, g, s)[x]; s \text{ is visited at root level} \\ \sum_{1 \leq i \leq N} A(i, s)[x] * dfs(t, g)[i]; \text{ otherwise} \end{array} \right\}$$

[0096] For the sake of simplicity, the learning system's ability to capture correlated activities was described in terms of a primary and secondary dataspace only. However, the invention is not limited to correlations between only two dataspace. The learning system can also capture user activities and preferences when the user enters a first dataspace, enters and performs an action in a second dataspace, and enters and performs yet further actions in a third dataspace. In fact, using the principles and techniques described above, the learning system can create $N!$ number of correlation signatures.

[0097] Signature Clustering

[0098] As described above, the users' activity and preferences are stored as vectors. Capturing the users' activity and preferences as vectors has many benefits. Among these is the fact that because signatures are vectors, one can easily determine when two signatures are nearly identical. For example, let $\tilde{x}$ and $\tilde{y}$ be two signatures in the same vector space. Lower case letters in bold will generally denote normalized probability vectors in appropriate dimension. A tilde over the top of a vector will generally denote unnormalized vectors. Greek symbols and lower case letters, without bold, will generally denote scalars. If

$$\text{vector angle}(\tilde{x}, \tilde{y}) = \frac{\tilde{x}^T \tilde{y}}{\|\tilde{x}\| \|\tilde{y}\|} \geq (1 - \epsilon), \quad (\text{Equation 7})$$

[0099] where ϵ is a small fraction in the vicinity of 0.01, then the two signatures are nearly identical. Equation 7 states that if the cosine of the angle between the two signatures is small enough, then they are nearly identical, up to a scaling factor. The scaling factor recognizes that two vectors may have different magnitudes, but still be overlapping. For example, a first user has a genre vector of {(sports 20); (world news 50)} and a second user has a genre vector of {(sports 60); (world news 150)} where the first value within the vector elements is the genre and the second value is minutes watched per day. Although the magnitudes of these two vectors are different, the genres and ratio of minutes of sports to world news is identical. Thus, the learning system identifies these two signatures as nearly identical. The learning system can exploit this aspect in a variety of ways, as described in greater detail below.

[0100] In one illustrative implementation, the system clusters signatures that are similar into subsets. The subsets can then be used for many purposes, such as promotions, targeted advertisements, etc. For example, if several users of a particular cluster have watched a certain television program, this television program will be recommended to other users of the cluster who have yet to view the program. Similarly, outside of the television context, if users of a particular cluster have purchased a given item, ads for this item are presented to the other users of the cluster who have not purchased the item.

[0101] The notion of viewing signatures as vectors can be exploited to determine the similarity between the signatures by using Equation 7. Each cluster represents nearly identical signatures. Initially, the procedure starts with singleton clusters, and recursively collapses them until no more merging is possible. An example of pseudo-code that generates clusters of signatures is provided below:

```

PROCEDURE
  Inputs:
    1. N signatures  $s_1, s_2, \dots, s_N$ 
    2. Tolerance threshold  $\epsilon$ ,  $0 \leq \epsilon \leq 1.0$ 
  Outputs:
    1. Sets  $\Psi_1, \Psi_2, \dots, \Psi_c$  containing the signature clusters
  BEGIN
    1. Initially define singleton sets  $\Omega_j := \{s_j\}; 1 \leq j \leq N$ 
    2. merged := FALSE
    3. for  $1 \leq i \leq N - 1$  do
      a. if set  $\Omega_i = \emptyset$ , continue
      b. for  $i+1 \leq j \leq N$  do
        i. if set  $\Omega_j = \emptyset$ , continue
        ii. If for every  $x \in \Omega_i$  and every  $y \in \Omega_j$ 
            vector_angle(x, y)  $\geq (1 - \epsilon)$ , then
            A.  $\Omega_i := \Omega_i \cup \Omega_j$ 
            B.  $\Omega_j := \emptyset$ 
            C. merged := TRUE
        end_if
      end_for
    end_for
    4. if merged = TRUE, go to step 2.
    5. c := 0
    6. for  $1 \leq i \leq N$  do
      a. if  $\Omega_i \neq \emptyset$ , then
        i. c := c + 1
        ii.  $\Psi_c := \Omega_i$ 
      end_if
    end_for
  END
END_PROCEDURE

```

[0102] Signature Decomposition

[0103] In addition, the learning system can decompose one or more signatures to approximate the number of family members in a signature, the gender of each family member, and the ages of the family members. The system uses a nonlinear optimization model to approximate these values. In this example, the learning system uses genre signatures; similar models apply to channel and microgenre signatures.

[0104] The technique starts with historical, human behavioral, statistical data on television viewing habits obtained from generally available data (such as the user viewing preferences available from the DVB Project). In particular, Δ is a set of all genres that are available to the viewing public. Upper case Greek letters in bold font generally denotes sets. Thus, a collection of probability distributions exists, namely

$$f_g^y(t) \equiv \text{Probability that a genre } g \text{ would be watched by a} \quad (\text{Equation 8})$$

person of gender y and age t ; $y = \{\text{male, female}\};$

$$0 \leq t \leq \infty; g \in \Delta$$

[0105] The learning system also provides a signature s that collectively describes the viewing habits of the household. The illustrative maximum likelihood estimation problem formulated below defines the most likely set of household

members that may produce the signature s . For the purposes of this example, all vectors, normalized or unnormalized, are lower case bold letters.

[0106] The inputs to the optimization problem are the probability distribution functions $f_g^y(t)$ and a family signature s that collectively represents the activities of a household. The outputs are n , the number of members in the household, where $1 \leq n$, the age and gender of each family member i , where $1 \leq i \leq n$, and a set of signatures $s_1, s_2, \dots, s_n$, where signature s_1 corresponds to family member i . Further, let $N = |\Delta|$, the cardinality of the Genre set, $\Phi = \text{set of all nonnegative basis matrices } B \text{ for the vector space } \mathbb{R}^{+N}$ (i.e., $B = [b_1, b_2, \dots, b_N]$, where b_1 is a nonnegative N -vector and $b_i, 1 \leq i \leq N$, are linearly independent, and for any vector $s \in \mathbb{R}^{+N}$,

$$s = \sum_{i=1}^N \alpha_i b_i,$$

with $\alpha_i \geq 0$).

[0107] The decision variables are as follows: basis matrix $B \in \Phi$, variables $x_1, x_2, \dots, x_N$, which represent the ages of family members with signatures corresponding to the basis vectors in B , and variables $y_1, y_2, \dots, y_N$, which represent the gender of the corresponding family members.

[0108] For the purpose of optimization, it is necessary to define an objective function to maximize. Towards this end, the system uses an intermediate function, as follows:

$$s = \sum_{i=1}^N \alpha_i b_i, \alpha_i \geq 0, \quad (\text{Equation 9})$$

for any vector $s \in \mathbb{R}^{+N}$ and any basis Matrix $B \in \Phi$

$$h(v, x, y) \triangleq \prod_{1 \leq k \leq N} f_k^y(x) v^{(k)}, \quad (\text{Equation 10})$$

where $v^{(k)}$ is the k^{th} component of vector v

[0109] Function $h(v, x, y)$ evaluates the likelihood probability of a person with age x and gender y having a signature v . Note that the system is taking the product of all the components of vector v . Thus, the maximum likelihood estimation becomes

$$\text{Maximize}_{B \in \Phi} \sum_{1 \leq j \leq N} h(\alpha_j b_j, x_j, y_j) \quad (\text{Equation 11})$$

[0110] subject to the following constraints:

$$s = \sum_{i=1}^N \alpha_i b_i, \alpha_i \geq 0; \quad (\text{Equation 12})$$

$$1 \leq x_j \leq \infty; 1 \leq j \leq N; \quad (\text{Equation 13})$$

$$y_j = \{0, 1\}; 1 \leq j \leq N; \quad (\text{Equation 14})$$

[0111] This optimization problem can be shown to be NP-Hard (see M. R. Garey, and D. S. Johnson, *Computers*

and Intractability A Guide to the theory of NP-completeness, W. H. Freeman and Company, New York, 1979, herein incorporated by reference), since any technique needs to search over the space of all bases in R^{+N} and the fact that the y variables are integer variable. This problem has some similarities to another nonstandard class of optimization problems known in the literature as semi-definite programs. An approximate solution to this problem can be achieved using an illustrative technique described below.

[0112] The estimation technique uses a first approximation by converting these discrete variables to continuous variables with bounds. This makes the transformed problem amenable to differentiable optimization. The technique also uses a second approximation by identifying a subset of the set of all bases Φ as the set of bases in R^{+N} that are permutations of the coordinate basis matrix and restricts the search to this set. Given any basis from this set, the inner iteration involves a steepest ascent technique in variables (α_j , x_j , and z_j) to obtain a local maximum (where z_j is a continuous approximation of y_j). The iterations are terminated when no improvement in the objective function occurs. After termination, the gender variables that are fractional are rounded/truncated using a rounding heuristic, described below. Given a fixed basis matrix, the transformed maximum likelihood estimation problem becomes a continuous maximum likelihood estimation problem and is given by the following equations:

$$\text{Maximize } \sum_{1 \leq j \leq N} h(\alpha_j b_j, x_j, z_j) \quad (\text{Equation 15})$$

[0113] subject to the following constraints:

$$s = \sum_{1 \leq j \leq N} \alpha_j b_j \quad (\text{Equation 16})$$

$$1 \leq x_j \leq \infty \quad 1 \leq j \leq N \quad (\text{Equation 17})$$

$$0 \leq z_j \leq 1 \quad 1 \leq j \leq N \quad (\text{Equation 18})$$

$$\alpha_j \geq 0 \quad 1 \leq j \leq N \quad (\text{Equation 19})$$

[0114] An example of pseudo code for solving the above continuous maximum likelihood estimation problem is given below. The pseudo code consists of an outer iteration and inner iteration. In the outer iteration, the code iterates through basis matrices. While in the inner iteration, the code employs the steepest-ascent optimization technique to obtain the optimal solution, given a basis matrix.

[0115] The steepest-ascent optimization technique has three steps. First, the optimization technique obtains the gradient of the objective function at the current iteration. This is done in step 2.c.ii, set forth below, using difference approximations. The technique of using difference approximations, as well as other numerical and matrix techniques can be found in D. G. Luenberger, Linear and Nonlinear Programming, second edition, Addison-Wesley publishing company, Reading Mass., 1989, herein incorporated by reference. Second, the optimization technique projects the gradient onto the null space of B, to obtain the ascent direction d (step 2.c.iii). Third, the optimization technique obtains the optimal step length along d (step 2.c.iv). In the

field of optimization, this is called a step-length computation and involves a one-dimensional optimization. The inner iterations proceed until no more improvement in the objective function is possible. After this, the basis matrix is changed and the inner iterations are reinitiated. Finally, rounding heuristics (such as those in G. L. Nemhauser, and L. A. Wolsey, *Integer and Combinatorial Optimization*, John Wiley & sons New York, 1988, herein incorporated by reference) are employed to round off the fractional variables.

[0116] In the pseudo code set forth below, I is an identity matrix of order N, and P_i is the i^{th} permutation matrix in the sequence of permutations of the index set $\{1, 2, 3, \dots, N\}$.

PROCEDURE

Inputs:

1. The historical probabilities $f_k^v(x)$
2. The family stochastic signature s

Outputs:

1. Number of family members
2. Sex of family members
3. Age of family members
4. Individual signatures of family members

BEGIN

1. Initialize:

- a. Current basis matrix $B := I$
- b. Iteration counter $i := 0$
- c. Permutation matrix $P := I$
- d. $\text{newOuterObj} := 0$; $\text{oldOuterObj} := -\infty$
- e. $\alpha_j := s_j \quad 1 \leq j \leq N$
- f. $x_j := 1 \quad 1 \leq j \leq N$
- g. $z_j := 0 \quad 1 \leq j \leq N$
- h. stopping tolerance for inner iteration $\epsilon := 0.001$
- i. stopping tolerance for outer iteration $\beta := 0.0001$
2. While $((\text{newOuterObj} - \text{oldOuterObj}) / |\text{oldOuterObj}| > \beta)$ Do
 - Do//outer iteration
 - a. $\text{oldOuterObj} := \text{newOuterObj}$
 - b. Initialize inner iteration; $\text{newInnerObj} := 0$; $\text{oldInnerObj} := -\infty$
 - c. while $((\text{newInnerObj} - \text{oldInnerObj}) / |\text{oldInnerObj}| > \epsilon)$ Do
 - i. $\text{oldInnerObj} := \text{newInnerObj}$
 - ii. Compute the gradient vector $g := [\partial h / \partial \alpha_j], [\partial h / \partial x_j], [\partial h / \partial z_j]$ using difference approximation.
 - iii. Project the gradient vector g on to the null space of matrix B to obtain the direction vector $d := g_{\perp B}$
 - iv. Compute the optimal step length δ along the direction d.

$$v \cdot \begin{bmatrix} \alpha_j \\ x_j \\ z_j \end{bmatrix} := \begin{bmatrix} \alpha_j \\ x_j \\ z_j \end{bmatrix} + \delta d$$

$$\text{vi. new InnerObj} := \sum_{1 \leq j \leq N} h(\alpha_j b_j, x_j, z_j)$$

end While

d. $i := i + 1$; set $P_i :=$ Next permutation matrix in the sequence

e. $B := P_i I P_i^T$

f. $\text{oldOuterObj} := \text{newOuterObj}$

g. $\text{newOuterObj} := \text{newInnerObj}$

endWhile

3. Use rounding heuristics to set fractional z_j variables to the nearest integer value to obtain variables y_j .

4. Compute $n :=$ number of α_j that are greater than 0.

5. Output the optimal solution:

h. Output n as the number of family members

i. For $1 \leq j \leq N$ Do

i. if $\alpha_j > 0$ then

a. output $\alpha_j b_j$ as the signature of person j

b. output x_j as the age of person j

c. output y_j as the sex of person j

EndIf

EndFor

END

END PROCEDURE

[0117] Signature Aging

[0118] The learning system also provides for remembering past behavior and integrating it into the current signature. The fundamental reason for remembering the past is to infer a recurring pattern in the user behavior, and use the inference to aid in future navigation. Exponential smoothing is a way of gradually forgetting the distant past, while giving more relevance to the immediate past. Thus, activities done yesterday are given more relevance than activities done two days ago, which in turn is given more importance than the activities done three days ago, and so on (see V. E. Benes, *Mathematical Theory of Connecting Networks and Telephone Traffic*, Academic Press, New York, 1965 for additional information, herein incorporated by reference). This technique has the added advantage of reducing the amount of computer memory consumed by the signature.

[0119] The learning system uses the concept of exponential smoothing in the context of learning the user's activities. For example, a set of activities for today is captured in a signature s , whereas all of the past activities are remembered in a signature s^* (s^* remembers all of the past, since the recurrence relation, given below, convolutes all of the past into a single signature). At the end of the day (when the present becomes the past), the system updates s^* by the recurrence relation

$$s^* = \alpha s + (1 - \alpha)s^*0 \leq \alpha \leq 1 \quad (\text{Equation 20})$$

[0120] In Equation 20, α is called the smoothing parameter and it controls how much of the past the system remembers, and how fast the past decays—larger the α , faster the decay. Expanding the above recurrence relation into a recurrence equation illustrates the machinery of exponential smoothing. Where $s^{*(n)}$ denotes the past signature after n days and $s^{(n)}$ represents the activities done during the n^{th} day. Equation 20 expands into

$$s^{*(n)} = s^{(n)} + (1 - \alpha)s^{*(n-1)} + (1 - \alpha)^2 s^{(n-2)} + \dots + (1 - \alpha)^{n-1} s^{(1)} + \frac{(1 - \alpha)^n}{\alpha} s^{(0)}. \quad (\text{Equation 21})$$

[0121] Because $\alpha \leq 1$, Equation 21 clearly shows that less weight is given to the activities of that past. In some embodiments, all signatures for each macro class (channel, genre, microgenre) are smoothed using the exponential smoothing technique. The determination of when to decay a particular signature is based on the dataspace of the signature and the nature of activities performed in the dataspace. For example, in the television dataspace, a decay period of one day is a suitable period because television shows typically reoccur on a daily basis. Whereas the decay period for the telephone dataspace would be longer so as to decay at a slower rate than the television dataspace. The decay parameter, or smoothing parameter α , can be selected to control the degree of decay of the past user behavior.

[0122] The learning system also uses an adaptive decay technique to integrate the past signature information with the most recent signature information. This adaptive decay technique is based on a hybrid chronological-activity based

smoothing and provides improved results over strict chronology-based aging when applied to signatures that take into account the user's geographic location. This technique enables the influence of past activities to decay over time, while still preserving the past information during the user's absence from a particular geographic location for a stretch of time. In general, the past signature will be decayed if (1) a new activity has occurred in the geographic location and (2) the elapsed time since the last signature decay event is greater than a threshold. In essence, the system freezes the signatures when no activity is happening in a given location, effectively stopping time for that location. When next an activity occurs in that location, the system smooths the signatures based on elapsed time.

[0123] If a traditional smoothing technique were employed to decay the memory of the past once per day, for example, the signature values may decay to zero if the user were absent from the geographic location for an extended period. Thus, upon returning to that particular location, the user would effectively have to "retrain" the system by rebuilding the signatures corresponding to that location. The adaptive decay technique avoids this problem.

[0124] An illustration of an implementation of the adaptive decay technique follows. As mentioned above, signature decay occurs for all signatures in coordinates (t, g, s) (i.e., time t , geographic location g , and dataspace s), only when there is a new activity in (t, g, s) . In addition, a minimum time must elapse before decaying takes place. To account for long elapsed times, the system uses the concept of epoch time. Epoch time is the absolute local time since a certain distant past. The concept of epoch time can be found in current-day operating systems (e.g., Linux and WinCE) that fix a reference point in the distant past and return the elapsed time since that reference point. For the example below, T is the epoch time when some activity x happens in (t, g, s) . Note that the coordinate t is an integer denoting the discretized time denoting the time-of-day or time slot, whereas T is an epoch time. For use in the Equation 22 below, $\beta(t, g, s)$ is the decay threshold for signatures, $r(t, g, s)$ is the last time, in epoch units, that signatures in (t, g, s) were decayed, and $e(t, g, s)$ is a vector capturing a newly performed user action (i.e., current signature) with a duration/count metric (explained above) in position x and zeros in all other positions. This technique also uses the smoothing parameter α as described above. Equation 22, shown below, is one implementations the adaptive decay technique.

(Equation 22)

$$ks(t, g, s) = \begin{cases} \alpha e(t, g, s)[x] + (1 - \alpha)ks(t, g, s); & \text{if } T > r(t, g, s) + \beta(t, g, s) \\ ks(t, g, s) + \frac{\alpha}{1 - \alpha} e(t, g, s)[x]; & \text{otherwise} \end{cases}$$

[0125] Under this implementation, the system decays the signature if the time interval since the last decay is greater than the decay interval; in this case, the system performs a convex combination of the past activity and present activity. If the last decay has occurred more recently than the decay interval, then the historic signature is combined with the current signature, with a multiplier $\alpha/(1 - \alpha)$ applied to the current signature. The technique of using this multiplier optimizes storage. Typically, when performing an exponen-

tial smoothing operation, the past is the period of time up to time $r(t, g, s)$, and the present is the period of time from time $r(t, g, s)$ to time T . Under a typical application, the new activity x would be stored in a temporary storage, $ts(t, g, s)$, along with all additional subsequent activities, until the time $r(t, g, s) + (\beta(t, g, s))$. At that time, the smoothing formula would combine the past with the new activities according to Equation 23.

$$ks(t, g, s) = \alpha ts(t, g, s) + (1 - \alpha)ks(t, g, s) \quad (\text{Equation 23})$$

[0126] The system avoids the need for temporary storage by combining each new activity with the past signature as each new activity occurs, using the multiplier described above to offset what would otherwise be a premature composition. This ensures true exponential smoothing. Although the above discussion involved only the keyword signatures, ks , the same principles and techniques apply to all other signatures described herein.

[0127] Use of Signatures to Personalize Content

[0128] As mentioned above, one illustrative use of the learning system is to enhance the user experience during a search procedure. In one illustrative implementation, the various individual, aggregate, and program signatures reside on a server system that contains a set of content items (e.g., television programs, television channels, movies, etc.). The server system uses the signatures to personalize search results provided to users of the system. In particular, the results obtained through a query are identified and reordered by promoting relevance values of individual search results based on the set of signatures. For example, in a system employing an incremental search method (as described in the above incorporated U.S. Patent Applications), the system begins searching for content item results as the user enters letters of a text query. The system identifies content items as candidates for presentation based on comparing the letters of the text query with descriptive terms associated with the content items. Each of these content items is associated with a base relevance value that measures the popularity of the item in the overall population. The system uses these base relevance values to rank which content items are likely sought by the user. Higher base relevance values indicate a higher overall popularity, thus, these items are assumed to be of more interest to the user than items with lower base relevance values.

[0129] However, as explained in greater detail below, the system modifies the base relevance values based on the set of user signatures. Thus, if the set of signatures indicates, given the particular time and day of the search, that it is likely the user is searching for a program with the genre of news, the system will promote the relevance values of programs with a genre of news that match the user's query text. Likewise, the system can use the channel and microgenre data associated with the content items in conjunction with the channel and microgenre signatures to promote the base relevance values. The final relevance weights of each item determine if the item is included in the result set and help determine its position in a list of results. Many different promotion techniques can be implemented; one example is the "ruthless promotion technique", described below.

[0130] The ruthless promotion technique ensures that any particular search result item that has a nonzero probability in

a user signature will have its relevance value boosted such that it will be higher than any other search result items having a zero probability value in the same user signature. For use in Equation 24 below, K is the number of search results retrieved with relevance numbers $r_1, r_2, \dots, r_K$, and M is the maximum value any relevance can have, based on the general popularity of the search result. Typically, search engines assign a relevance number to query results based on ranks with some maximum bound. These ranks can be a measure of the popularity or relevance of the items based on popular opinion. Search results are displayed in the shelf space, sorted in descending order of relevance based on these ranks. (Herein, the phrase "shelf space" refers to the portion of a display screen of a device that displays the search results in response to a user query. This portion can be organized as a column of text boxes in some implementations.) The values $p_1^{(1)}, p_2^{(1)}, \dots, p_K^{(1)}$ are the channel signature probabilities ($0 \leq p_i^{(1)} \leq 1$) assigned by the learning system (typically, most of the $p_i^{(1)}$ will be 0). The superscripts on the probabilities refer to type of signature, e.g., channel, genre, or microgenre. The ruthless promotion technique computes new relevance numbers $\tilde{r}_1, \tilde{r}_2, \dots, \tilde{r}_K$ as

$$\tilde{r}_i = \begin{cases} \lfloor (M+1)e^{p_i^{(1)}} \rfloor; & p_i^{(1)} > 0 \\ r_i; & p_i^{(1)} = 0 \end{cases} \quad (\text{Equation 24})$$

[0131] The search items are then reordered using the new relevance numbers. For example, a user had watched the channels "CARTOON NETWORK" and "COMEDY CHANNEL" in the past, with the signature probabilities 0.7 and 0.3 respectively. The generic relevance numbers for channels, based on popular opinion, are 500, 300, 100, and 70, for "CBS", "CNBC", "COMEDY CHANNEL", and "CARTOON NETWORK", respectively with a maximum bound of 1000. Table 1 and Table 2 show the displayed results and their corresponding relevance values, when a query character "C" is typed. Table 1 shows the order of the query results without the benefit of the learning system, and Table 2 shows the order of the results using the ruthless promotion technique. As can be seen, the user convenience is enhanced, because fewer scrolls are required to access the most likely watched channels.

TABLE 1

Channel	Relevance Number
CBS	500
CNBC	300
COMEDY CHANNEL	100
CARTOON NETWORK	70
...	...

TABLE 2

Channel	Relevance Number
CARTOON NETWORK	2015
COMEDY CHANNEL	1351
CBS	500
CNBC	300
...	...

[0132] Other promotion techniques are within the scope of the invention, including, for example, techniques that do not necessarily ensure that a search result item that has a nonzero probability in a user signature will have its relevance value boosted such that it will be higher than any other search result items having a zero probability value in the same user signature. In particular, because there are six signatures capturing the user activity at any time of day—channel, genre, and microgenre for given time slot, and their corresponding composite signatures, these signatures are combined to compute new relevance weights. Equation 24 above shows the use of channel signature for promotion. In the example below, there is an inherent importance in these signatures, from more refined to more coarse. This variant of the ruthless promotion technique considers an aggregated promotion formula, as follows:

$$\tilde{r}_i = \left[\sum_{1 \leq k \leq 6} (M+1)^{k e p_i^{(k)}} \right] \quad (\text{Equation 25})$$

[0133] In Equation 25, the superscript on the probabilities, $p_i^{(k)}$ refers to time slot channel, microgenre, genre, followed by composite channel, microgenre, and genre signatures, with increasing values of k , respectively. Since they are being multiplied by powers of $(M+1)$, a natural relevance importance is implied.

[0134] Signatures corresponding to overlapping periodicities are also combined to provide an aggregate signature for a particular time slot. The probabilities in the aggregate signature can be used with the promotion techniques above to identify and rank the results of search. In order to form an aggregate signature, the vectors from the overlapping signatures are added together and renormalized. For example, for a particular time slot, a user has a first signature with a periodicity of every Monday and a second signature with a periodicity of every weekday. The first signature is the genre vector $\{(0.2 \text{ news}), (0.8 \text{ sports})\}$; the second signature is the genre vector $\{(0.1 \text{ comedy}), (0.4 \text{ news}), (0.5 \text{ sports})\}$. To form an aggregate signature, the system first arithmetically combines the two vectors to produce the new vector $\{(0.1 \text{ comedy}), (0.6 \text{ news}), (1.3 \text{ sports})\}$, and the system then normalizes the new vector by dividing each numerical element by the sum of the numerical elements of the vector, i.e., 2.0. Thus, the aggregate, normalized genre probability vector of the two overlapping signatures is $\{(0.05 \text{ comedy}), (0.3 \text{ news}), (0.65 \text{ sports})\}$.

[0135] Seminormalization of Signatures

[0136] In one implementation, the learning system uses integer arithmetic for relevance promotion computations. In particular, all probabilities are represented as integers, appropriately scaled. One compelling motivation for using integer arithmetic is to make the learning system portable to disparate hardware and operating systems, some of which may lack floating arithmetic.

[0137] The learning system uses a seminormalization approach to weight more recent activities more heavily in the signatures, while deemphasizing, but still retaining, the information from more distant activities in the signature. Thus, when personalizing services or content provided to the user, the system is more heavily influenced by more recent activities. The basic idea of this seminormalization approach is to make the long-term memory coarse-grained by buck-

eting small probabilities that result from less common user preferences and/or preferences captured in the more distance past, while still bounding the range of values by using a small relevance scale factor. This approach allows the system to capture both small and large probabilities without requiring a large dynamic range to define the probability bounds. Thus, a small scaling factor is used to distinguish between the relatively more probable activities in the captured in the signatures, while the relatively less probable activities are not lost due to truncation errors.

[0138] An illustrative technique for converting an unnormalized signature, x , into a seminormalized signature is provided below. In signature x , all elements x_i are nonnegative integers, representing the activity intensity. Signature x is an N -dimensional vector, and x has at least one positive element. The value d is the infinity norm of x and

$$\|x\|_{\infty} \triangleq \sum_{1 \leq i \leq N} x_i;$$

thus, d is the normalizing sum. The vector p is the normalized probability vector corresponding to x ; therefore,

$$p_i = \frac{x_i}{d}; 1 \leq i \leq N.$$

[0139] In order to seminormalize the signature x , the system uses a fine-grain memory threshold of probabilities, K , represented in log scale; i.e., K is a positive integer denoting a probability threshold 10^{-K} . All probabilities $\geq 10^{-K}$ will be scaled in fine-grain, and all probabilities between 0 and 10^{-K} will be scaled in coarse-grain with bucketing. The system also uses a positive integer, S , as the seminormalization range represented in log scale. After seminormalization, a probability value of 1 is scaled to 10^S . The largest value in the seminormalized vector would be 10^S . Although not required, S can be equal to K . For use in the equations below, let $t=10^{-K}$, $1=K+2$, and $u=10^S$. Finally, y is the seminormalized vector corresponding to p . Thus, $y=f(p, K, S)$, where f is the function implementing the seminormalization algorithm. The function f is not an invertible function.

[0140] Each element i of the seminormalized vector y is defined by the Equation 26.

(Equation 26)

$$y_i = \begin{cases} 1; & 0 \leq p_i 10^{-2K} \\ v+2; & 10^{(-2K+v)} \leq p_i < 10^{(-2K+v+1)} \\ & 0 \leq v \leq K-1 \\ \left(\frac{(10^K-1-u)}{(10^K-1)} + \frac{10^K(u-1)p_i}{(10^K-1)} \right); & 10^{-K} \leq p_i \leq 1 \end{cases}$$

[0141] The first 2 parts of Equation 26 define the coarse-grain bucketing, and the last part of the equation defines the fine-grain scaling. FIG. 7 shows a pictorial representation of Equation 26. The X-axis is shown in log scale. In FIG. 7, $S=K$, and there are K buckets of width 0.1. The buckets start with the bucket having a left boundary 10^{-2K} and ending

with the bucket with the right boundary 10^{-K} . There is a special underflow bucket for any probability $<10^{-2K}$. Each p_i falling within a bucket is scaled to a fixed count. For probabilities larger than 10^{-K} , the p_i is scaled using a linear equation. The slope of the linear scaling equation in plot is approximately 10^S with the intercept at $(K+2)$.

[0142] An example of this approach as applied to an electronic phonebook application on a mobile phone is provided below. In this example, each outgoing call is counted and stored in a raw signature. The system scales the call counts by a large number so that truncation errors in the smoothing and aging process, due to integer arithmetic, are reduced. FIG. 8 illustrates a raw phonebook signature **800** with 6 entries. The row names in signature **800** indicate the person called, and the values are the scaled frequency, after aging and smoothing. Thus, the value 1 represents the long-term memory of a phone call made to John, perhaps many years ago, and not repeated again. Similarly, the entry corresponding to Jane has a signature value of 5. This value can be interpreted two ways: (1) Jane was called as long ago as John, but with a frequency five times greater than John; or (2) Jane was called more recently with the same frequency as John. The larger values represent short-term memories of calls made in the recent past. The normalized probabilities of these events are shown in a probability vector **810**.

[0143] It is clear the dynamic range of probability vector **810** is quite large. Using the techniques described above, with $K=S=4$, the system generated a seminormalized vector **820**. The system has collapsed the memory of John and Jane into the underflow bucket, thus making them indistinguishable. Some differentiation has been made for Debbie and Naren, although these entries also represent long-term memories, and therefore, the exact difference in frequency or recency of calls is not retained. However, the system captures the precise relative values of Simon and Marty. The values in the seminormalized vector are completely bounded and suitable for stable relevance promotion.

[0144] Activity Spike Detection

[0145] The learning system is also useful for detecting sudden bursts of activity (i.e., spike detection) at a particular time and day for a particular search item, e.g., a certain television channel. The system can use these spikes of activity to temporarily boost the base relevance values of particular items that have sudden popularity. Typically, spikes happen when an emergency or crisis has happened in the recent past. For example, if a particular news channel is presenting a late-breaking news story that is attracting a high number of viewers, the system will recognize the sudden popularity of the news program and boost its base relevance in recognition of the fact that other users of the system may also be interested in the news story.

[0146] In general, when the collective activity level associated with a particular content item is above a certain threshold attributable to statistical variations, then the activity level is considered a spike. The learning system analyzes the current and past activity levels by collectively examining all of the signatures of the user population. If each user is considered an independent random variable whose probability of watching a program is encoded in a stochastic signature, then the collection of all these independent random variables provides a measure of the overall popularity of the content item. Thus, the system employs these signatures to derive a joint probability distribution of the number of users

watching a given program at a given time. Thus a new type of signature, herein a “program signature”, $r_k^{(i,t)}$, is defined in Equation 27.

$$r_k^{(i,t)} := \text{Probability that a program } i \text{ is being} \quad (\text{Equation 27})$$

watched by k users at time t ; $0 \leq k \leq N$

[0147] An example of a technique for obtaining the program signature is provided below. In general, when the activity level associated with a particular content item exceeds a certain inherent randomness value predicted by the program signature, the system identifies such activity as a spike.

[0148] The system creates a set of program signatures, each of which is a statistical convolution of all individual signatures in the population that have watched the particular program. By convolving the individual signatures, the system creates an aggregate mean and standard deviation of the activity level associated with the given program. Thus, a program signature captures the fraction of all of the current users interacting with the system that are currently watching the given program. Because the number of users interacting with the system changes over time, the fraction of users watching a particular program changes over time as well. The system captures this information by creating program signatures for the various time slots.

[0149] These signatures estimate the mean and higher moments of the probability distribution of the number of people accessing this program in terms of fractional probabilities. The aggregate signature and related statistical measures define the normal level of activity for the particular search item. Thus, by continually monitoring the current level of activity for the particular search item at a given time, the system can detect if the current activity is above or below the normal level. If the activity exceeds a certain threshold, the system adjusts the reordering technique to temporarily boost the relevance of the particular search item to recognize that the item may be of particular interest.

[0150] An example of creating an aggregate signature is provided below. For the sake of simplicity, the example is restricted to one day, one time slot, one content item, i , and a single category of signature (e.g., channel, genre, or microgenre). This technique for finding the aggregate signature is applied to all time periods, all days, all search items, and all signatures. In the following example, N is the number of users using the system, $q_i^{(j)}$ is the normalized signature value of user j for item i (i.e., the fraction of time user j watched program i) where $1 \leq j \leq N$, Ψ is the index set $\{1, 2, \dots, N\}$, Φ_m is the set of subsets of Ψ of size m where $0 \leq m \leq N$, and X is a random variable denoting the number of users currently watching program i .

[0151] The unnormalized probability that there are m users watching program i , herein r_m , is determined by Equation 28.

$$r_m = \sum_{\Theta \in \Phi_m} \prod_{1 \leq k \leq m} q_i^{(j_k)}, \text{ where } \Theta = \{j_1, j_2, \dots, j_m\} \quad (\text{Equation 28})$$

[0152] The normalization constant, G , is given by Equation 29.

$$G = \sum_{0 \leq m \leq N} r_m \quad (\text{Equation 29})$$

[0153] The probability density function of X , $f_X(m)$, the mean of X , μ_X , and the standard deviation of X , σ_X are now given by the following equations:

$$f_X(m) = \frac{1}{G} r_m; 0 \leq m \leq N \quad (\text{Equation 30})$$

$$\mu_X = \sum_{0 \leq m \leq N} m f_X(m) \quad (\text{Equation 31})$$

$$\sigma_X = \sqrt{\sum_{0 \leq m \leq N} (m - \mu_X)^2 f_X(m)} \quad (\text{Equation 32})$$

[0154] The system monitors the number of users watching program i . Chebychev's inequality dictates that, with 96% confidence, the random variable X cannot be above $\mu + 5\sigma$ due to inherent randomness. Thus, whenever the number of users watching program i goes beyond $\mu_X + 5\sigma_X$, the system identifies this as a spike of activity. The system can temporarily boost the base relevance of program i in queries for which program i is a candidate in recognition of the fact that the user may be interested in the same program. The relevance can be boosted by a predetermined amount, or it may be boosted by an amount that is proportional to the increase in viewing of the program. In addition, the system can use a variety of multiples of σ_X (not only $5\sigma_X$) to determine when a spike of activity is occurring.

[0155] The system can also infer the overall relevance of particular search items using the aggregate signatures. As described above, the system computes the mean of the statistical convolution of N signatures, N being the number of system users. Using this mean value, the system generically reorders the search results even in the absence of a signature for a particular user. Thus, the user benefits from the system's knowledge of the popular opinion of various search items, and these popular opinions are used to identify and order search results for presentation to the user. For example, if the aggregate signature has a large mean for the television program "The Apprentice", then any user who does not have a personal signature will have this item in the top shelf on an appropriate query (the query, for instance, can be "trump", which is a microgenre of the program "The Apprentice").

[0156] FIG. 9 illustrates an example of detecting an increased level of activity associated with a content item (i.e., an activity spike). A normal level of activity 905, as determined using the techniques described above is shown. Normal level of activity 905 is based on the aggregate signatures. As the system is being used, a current level of activity 910 is generated using continuously calculated aggregate signatures based on the current content items usage or activity. Upon detecting an increase in activity level 915, which is beyond a specified threshold, the system identifies the content item as having a spike of activity, and the system promotes the ranking of that content item, as described above.

[0157] The learning system also allows accessing rare search items using preprocessing. In some implementations described above, the search engines work by first gathering significant amounts of results matching the query, and filtering out low relevance results before applying a promotion technique. This technique has several advantages, including increasing the speed of the system and reduces network bandwidth required. However, a specific user may be interested in an item having low overall popularity that is filtered out of the results before applying a promotion technique. In the absence of a signature, this rare item may never be presented in the search results (this rare item is sometimes referred to as the "long tail" in the probability distribution sense).

[0158] In order to capture the rare item in the ordered search results, some implementations of the system compute the relevance before filtering, using the promotion techniques described above or other promotion techniques. Thus, the rare item is ranked highly for the particular user, allowing him or her to access the item with ease. Here, signatures enable fine-grain customization and increase user satisfaction.

[0159] An inherent feature of the stochastic signature mechanism is the probabilistic nature of the signature entries, i.e., the signature entries are all normalized probabilities. This enables the system to export the signatures to other, potentially unrelated systems, with ease. For example, over some period of time, the television system interface described above learns that, in general, a given user prefers the Spirituality genre 50% of the time, Sports genre 40% of the time, and the Seinfeld show 10% of the time. In response, the system creates a set of signatures for the user that captures these preferences. The user can elect to share this signature information with other systems.

[0160] Therefore, when the user registers with a website that sells books, the user can elect to share his signature information with this website. Because the signature information is stored in terms of normalized probabilities, the signature can be easily imported into the website that is configured to utilize such probability information. In addition, the website need not have an identical set of genres as that of the television system in order to use the signature information. For example, the website may not have "Seinfeld" defined as a genre or category of books. In this case, the website can simply renormalize the signature by removing the irrelevant entries, i.e., Seinfeld, and determining new normalized probabilities for the remaining genres. Thus, the new normalized probabilities for the user would be 56% for Spirituality and 44% for Sports. Sharing signatures in this way obviates the need for relearning in the new system. Also, different subsets of signatures can be shared for different systems.

[0161] Signature Based Preference Service

[0162] As explained above, the learning system captures the user's preferences across multiple dataspace. In addition, portions of the learning system can be incorporated into various user client devices, thereby enabling the system to capture the user's preferences across multiple devices. For example, the system can track the user's actions performed on a mobile telephone, a television system, a handheld computer device, and/or a personal computer. This enables the system to provide personalized services to the user across the multiple dataspace and multiple devices. Thus,

user preferences expressed on one device can be used to personalize the user interactions on a different device.

[0163] Likewise, the learning system can provide the learned user preferences captured in the various signatures to third-party service providers. The information provided to third-party service providers allows the service providers to personalize the services for the user outside of the learning system. In such an implementation, the learning system determines what preference information to provide to the service providers based on the nature of the services provided. The learning system can provide this information on a per transaction basis, or the system can periodically synchronize a set of signatures stored by the third-party service provider. Furthermore, the user can configure the learning system so as to control which third-party service receives user preference information.

[0164] By providing a centralized system that learns and stores the user's preferences, the learning system enables the user to avoid creating disconnected pockets of personalization associated with only one particular dataspace or device. Therefore, a user may immediately leverage the preference information contained in the user's signatures when interacting with a new service rather than having to wait for the new service to learn the user preferences. Thus, the learning system can provide personalization information to the third-party service provider to improve the user's experience with the service.

[0165] This comprehensive personalization across diverse user activities and devices is especially helpful to the user when the user interacts with the same service provider using different interface devices. Not only does the learning system capture the user's preferences from these diverse interactions, but the system also stores the details of the user's transaction for later retrieval by the user. For example, a user can book a flight through a travel website using a personal computer. The learning system captures the detailed information associated with transaction, e.g., the date of travel, the time of the flight, and the departure and destination city. At a later time, the user wishes to modify the travel reservations, and elects to do so using a mobile telephone. Because the system monitors the user's interactions with various service providers, the system recognizes that the user has placed a telephone call to the travel service. In response, the learning system automatically presents the user's upcoming travel itineraries on the user's mobile telephone or sends the information to the travel service's customer service agent with the user's consent.

[0166] In the alternative, if the user is presented with an automated voice response system, the learning system can send the relevant itineraries to the travel service (e.g., via an SMS message dispatched to the telephone number called or DTMF tones at the beginning of the telephone call), which would provide the travel service with a background context of the telephone call to improve the automated voice response system's response to the user voice commands. The power of a comprehensive personalization across diverse user activities and devices becomes very evident in voice based navigation applications. Comprehensive personalization can provide the necessary context that can vastly improve ambiguities in user input that plague these types of systems today.

[0167] FIG. 10 illustrates a part of the learning system for providing a context specific personal preference information service. In a preferred embodiment, a user device **1001a-c**

solicits a service, on behalf of the user, from a service provider **1002**. This can include, for example, making a telephone call to modify a travel itinerary or accessing a search engine to find some information. The context-sensitive personal preference information service **1003** enables the external service provider **1002** to provide a targeted response to the user based on user's prior activity, data access history, and the learned user preferences.

[0168] Service provider **1002** can also serve as the source of information and relevance updates for user device **1001a-c**. A network **1002** functions as the distribution framework and can be a combination of wired and wireless connections. The navigation devices can have a wide range of interface capabilities and include such devices as a personal or laptop computer **1001a**, a hand-held device **1001b** (e.g. phone, PDA, or a music/video playback device) with limited display size and an overloaded or small QWERTY keypad, and a television remote control system **1001c**, wherein the remote control has an overloaded or small QWERTY keypad. The navigation devices provide user activity data to the learning system via personal preference information service **1003** to create the various signatures. As mentioned above, in alternate embodiments, the user device can create the various signatures, and the signatures can be kept on the device. This enables the device to locally filter and order content items received from service provider **1002** and/or content items that reside on the device itself.

[0169] As described above, the learning system captures the user's preferences from the user's interactions with various dataspace. FIG. 11 illustrates the local tracking and strengthening of the personal preference signatures based on user activity and the content on a mobile device. For example, user interaction with a telephone book **1101**, media applications **1102**, email/calendar **1103**, and web browser **1104** are tracked, as well as when and where the interaction takes place. In addition to the user's interaction with these applications, the content that is coupled with these applications such as call logs **1101A**, music files **1102A**, email data/calendar appointments **1103A**, and browser cookies **1104A** are also tracked to capture the user's preferences. Aggregated actions and various signatures **1105** are captured by the learning system as described above.

[0170] The aggregated data and signatures **1105** are used by a wide variety of services, ranging from a local data prefetching service, in order to improve search performance, to a commercial third-party service provider, in order to target the user for a specific product offering. The sets of signatures generated by the learning system form an onion-like layered structure; the inner layers are specific and capture the exact details of the user's actions, while the outer layers characterize the user's general preferences. For example, the inner layers capture (1) the time and the location where the user performed an action, (2) the nature of the action (e.g. tuning to a channel or the purchase of a book, DVD, or airline ticket), and (3) the details of the action (e.g. the channel and/or program the user tuned to, the title of book the user ordered, or the departure and destination airports of an airline ticket purchase). This layered structure coincides with the various signatures created by the learning system. The inner layers correspond to the microgenre and keyword signatures, while the outer layers correspond to the genre signatures.

[0171] The service provider requesting the user's signature information can designate the degree of specificity of

user preferences desired. For example, a video content search engine wishing to use the user's signatures to order the results of a query may request specific information on which channels or program the user watched. A bookstore, on the other hand, may request broad user preferences of book tastes. The personal signature information sent in the later case would not be the individual instances of purchased books, but rather the broad user preferences at a genre level.

[0172] FIG. 12 illustrates the information flow when a user device 1203 makes a request to a service provider 1201 (step 1). The request contains a unique identifier that identifies the user or the user device. The identity could be an anonymous yet unique identifier of the device. For example, a one-way hash function of the device hardware identifier may be used to uniquely identify the device; there would be no way to reverse map to the actual device that generated the response, given the one-way nature of the hash function. In this case, the personal preference service 1202 has only have a set of unique device identifiers that share signatures for each user; there would be no identity information beyond the identifiers. In this way, the user's identity is maintained anonymous, yet responses matching user's preferences can be delivered.

[0173] In addition to the substance of the request, the communication from user device 1203 to service provider 1201 contains information that describes the context of the request, as explained below. Service provider 1201 communicates the substance of the request and the additional context information to personal preference service 1202 (step 2). The context information includes the identifier of the user device currently being employed, the location of the user device, if available, the time of the request, and general description of the action the user is performing (e.g., the fact the user is currently using a telephone versus playing media). The additional context information enables personal preference service 1202 to provide context-sensitive personal preference information to service provider 1201. Descriptive tags are assigned to the various actions the user can perform using the system. The system associates these descriptive tags with the signatures that are generated by the corresponding actions. In this way, personal preference service 1202 sends relevant preference information based on the tags sent by user device 1203 to service provider 1201 (step 3).

[0174] The relevant personal preference information is used by the service provider 1201 to send a targeted response to the user device 1203 (step 4). Additionally, service provider 1201 sends feedback to personal preference service 1202 about the targeted response that was sent (step 5). This feedback is used by personal preference service 1202 to adjust the personal actions signature of the user.

[0175] By disaggregating personal preferences through a standalone entity, i.e. personal preference service 1202, multiple service providers that provide different services can all benefit from the aggregated personal preference accrued across different service providers, different user actions, and different user devices. The end user gains immensely due to the key benefit of having targeted responses to many different types of requests. For example, a user who purchases books from Amazon.com gets the benefit of a targeted response when he goes to the Barnes & Nobles site using the techniques described above.

[0176] As described above, personal preference service 1202 can also be a centralized aggregator of user actions

across different devices. Thus, user actions performed on different devices, e.g., a mobile computing device, a home television with set-top box, and a personal computer, could all be aggregated to provide user preferences for identifying and ordering search results. For example, a user could initiate a remote recording for a favorite program using a mobile device, where the discovery of the desired program can be made easy by leveraging the user's viewing behavior on a television system. Thus, the available episodes of Seinfeld could be automatically displayed in the mobile device, for example, based on the fact that the user has viewed Seinfeld many times in the past on the television system.

[0177] FIG. 13 illustrates the information flow when a user device 1302 makes a request to a service provider 1301. In this scenario, the context sensitive personal preference information is sent along with the request (step 1) to generate a response (step 2). The personal preference data 1303 is locally resident on user device 1302. Additionally, personal preference data 1303 is updated (step 3) based on the response received from service provider 1301.

[0178] In another implementation of the learning system, a user device can serve as the personal preference provider in a peer-to-peer fashion for other user devices. For example, in a home entertainment network with more than one DVR (Digital Video Recorder), one DVR can serve as the personal preference provider for another DVR resident in the same home. When the user performs a search for content on a particular DVR, the other DVR in the home provides a personalized preference service to enable the user to find the desired content more quickly by leveraging the prior user viewing habits across the different DVRs.

[0179] In addition, a particular user can elect to share his or her signatures with another user. This can be accomplished in a peer-to-peer fashion as described above. In this case, the preferences learned for one user are used to personalize content results for another user. For example, the system will generate a set of signatures for a first user while that user selected various content from a book dataspace. These signatures encode the book reading preferences of the first user. A second user has a similar interest to the first user, and the second user wishes to select books related to similar topics as the first user. In this case, the first user can share his signature with the second user. The system then uses the first user's signatures to personalize the content results for the second user. In this way, the system enables the second user to benefit from the learned preferences of the first user without the second user having to train the system.

[0180] FIG. 14 illustrates different services 1401, for example, travel services (airline, car, and hotel), food services, entertainment services, and search engines services, that benefit from the context-sensitive personal preference service 1402. Although each service provider may have its own personalized services, when users first identify themselves, the services have no knowledge of the first time customer. The techniques disclosed herein increase the likelihood of the acquiring and retaining first time customers by offering targeted services immediately upon using the service. The techniques disclosed also enhance the first-time user experience. In contrast, without these techniques, users would have to create an account with a service and build an action history with that service before receiving personalized services.

[0181] Using the techniques described above, a user, for example, can go to any travel site and the site, without knowing the user and without requiring him to create an account or log in, can still offer the user personalized services based on the history of prior travel actions the user took on other platforms or web sites. Additionally, for services where comprehensive personalization is not in place, these services can leverage the personal preference service discussed above.

[0182] Because the learning system and personal preference service operate across multiple dataspace and multiple user devices, the user device configuration can vary greatly. FIG. 15 illustrates possible user device configurations for use with the learning system and the context-sensitive personal preference service. In one configuration, a user device 1509 can have multiple output capabilities, for example, a display 1501 and/or a voice output 1502. In addition, user device can have a processor 1503, volatile memory 1504, a text input interface 1505, and/or voice input 1506. Furthermore, user device 1509 can have remote connectivity 1507 to a server through a network and can have persistent storage 1508.

[0183] In another user device configuration, user device 1509 may not have local persistent storage 1508. In such a scenario, user device 1509 would have remote connectivity 1507 to submit the user's request to a server and retrieve responses from the server. In yet another configuration of user device 1509, the device may not have remote connectivity 1507. In such case, the learning system, personalization database, and signatures are locally resident on local persistent storage 1508. Persistent storage 1508 can be a removable storage element, such as SD, SmartMedia, or a CompactFlash card. In a configuration of user device 1509 with remote connectivity 1507 and persistent storage 1508, user device 1509 can use remote connectivity 1507 for a personalization data update or for the case where the personalization database is distributed between local persistent storage 1508 and a server.

[0184] It will be appreciated that the scope of the present invention is not limited to the above-described embodiments, but rather is defined by the appended claims, and these claims will encompass modifications of and improvements to what has been described. For example, embodiments have been described in terms of a television content system. However, embodiments of the invention can be implemented on a mobile phone to assist the user in retrieving personal contact information for individuals.

What is claimed is:

1.-22. (canceled)

23. A computer-implemented method comprising:

receiving, at a server and from a client device, a user interface interaction with a particular content item of a plurality of content items, wherein each content item of a subset of the plurality of content items is associated with:

a respective macro-theme that broadly characterizes the respective content item,

a respective micro-theme that, with respect to the respective macro-theme, narrowly characterizes the respective content item, and

a respective relevance value that indicates how relevant the respective content item is to a user profile;

identifying, based on the user profile, at least one content item of the subset of the plurality of content items,

wherein a micro-theme associated with the at least one identified content item comprises the micro-theme of the particular content item;

increasing a relevance value associated with the at least one identified content item of the subset, based on the user interface interaction;

ranking the plurality of content items in an order based at least in part on the respective relevance value associated with each content item of the subset of content items; and

generating for display the plurality of content items based on the ranking.

24. The method of claim 23, wherein ranking the plurality of content items is further based at least in part on criteria other than the respective relevance values of the subset of the plurality of content items.

25. The method of claim 23, wherein ranking the plurality of content items is further based at least in part on at least one of date, day, or time.

26. The method of claim 23, wherein ranking the plurality of content items is further performed based at least in part on user interface interactions with other user profiles.

27. The method of claim 23, wherein ranking the plurality of content items is further performed based at least in part on historical data indicating content items the user profile previously interacted with via an application.

28. The method of claim 23, wherein ranking the plurality of content items is further performed based at least in part on a stochastic signature comprising statistical values associated with a plurality of micro-themes, and wherein at least one micro-theme of the plurality of micro-themes is associated with a statistical value indicating its relevance to the user profile.

29. The method of claim 23, wherein a trained model is used to:

identify the micro-theme of the particular content item associated with the user interface interaction;

perform the identifying of the at least one content item of the subset of the plurality of content items; and

perform the ranking of the plurality of content items.

30. The method of claim 23, wherein the order comprises the plurality of content items sorted based at least in part on a descending order of relevance, wherein a content item having a higher relevance value has a greater probability of being positioned at a top portion of a display or other prominent portion of the display.

31. The method of claim 23, wherein:

the user interface interaction is further associated with a first context comprising at least one of a geographic location, client device identifier, date, day, or time associated with the user interface interaction, and the increasing further comprises:

comparing the first context with a second context associated with at least one content item of the plurality of content items;

determining, based on the comparing, that the second context associated with the at least one content item comprises an attribute of the first context; and

increasing the relevance value associated with the at least one content item.

32. The method of claim 23, wherein at least one content item of the plurality of content items comprises at least one of an audio item, a video item, a web content item, an advertisement, or any combination thereof.

33. A computer-implemented system, comprising:
 a database for storing metadata for a plurality of content items;
 a server comprising processing circuitry configured to:
 receive, at a server and from a client device, a user interface interaction with a particular content item of a plurality of content items, wherein each content item of a subset of the plurality of content items is associated with:
 a respective macro-theme that broadly characterizes the respective content item,
 a respective micro-theme that, with respect to the respective macro-theme, narrowly characterizes the respective content item, and
 a respective relevance value that indicates how relevant the respective content item is to a user profile;
 identify, based on the user profile, at least one content item of the subset of the plurality of content items, wherein a micro-theme associated with the at least one identified content item comprises the micro-theme of the particular content item;
 increase a relevance value associated with the at least one identified content item of the subset, based on the user interface interaction;
 rank the plurality of content items in an order based at least in part on the respective relevance value associated with each content item of the subset of content items; and
 generate for display the plurality of content items based on the ranking.

34. The system of claim 33, wherein the processing circuitry is configured to rank the plurality of content items further based at least in part on criteria other than the respective relevance values of the subset of the plurality of content items.

35. The system of claim 33, wherein the processing circuitry is configured to rank the plurality of content items further based at least in part on at least one of date, day, or time.

36. The system of claim 33, wherein the processing circuitry is configured to rank the plurality of content items based at least in part on user interface interactions with other user profiles.

37. The system of claim 33, wherein the processing circuitry is configured to rank the plurality of content items further based at least in part on historical data indicating content items the user profile previously interacted with via an application.

38. The system of claim 33, wherein the processing circuitry is configured to rank the plurality of content items further based at least in part on a stochastic signature comprising statistical values associated with a plurality of micro-themes, and wherein at least one micro-theme of the plurality of micro-themes is associated with a statistical value indicating its relevance to the user profile.

39. The system of claim 33, wherein the processing circuitry is configured to use a trained model to:
 identify the micro-theme of the particular content item associated with the user interface interaction;
 perform the identifying of the at least one content item of the subset of the plurality of content items; and
 perform the ranking of the plurality of content items.

40. The system of claim 33, wherein the order comprises the plurality of content items sorted based at least in part on a descending order of relevance, wherein a content item having a higher relevance value has a greater probability of being positioned at a top portion of a display or other prominent portion of the display.

41. The system of claim 33, wherein:

the user interface interaction is further associated with a first context comprising at least one of a geographic location, client device identifier, date, day, or time associated with the user interface interaction, and the processing circuitry is configured to perform the increasing by:

comparing the first context with a second context associated with at least one content item of the plurality of content items;

determining, based on the comparing, that the second context associated with the at least one content item comprises an attribute of the first context; and

increasing the relevance value associated with the at least one content item.

42. The system of claim 33, wherein at least one content item of the plurality of content items comprises at least one of an audio item, a video item, a web content item, an advertisement, or any combination thereof.

43. One or more non-transitory computer-readable media storing instructions which, when executed by one or more processors, cause performance of:

receiving, at a server and from a client device, a user interface interaction with a particular content item of a plurality of content items, wherein each content item of a subset of the plurality of content items is associated with:

a respective macro-theme that broadly characterizes the respective content item,

a respective micro-theme that, with respect to the respective macro-theme, narrowly characterizes the respective content item, and

a respective relevance value that indicates how relevant the respective content item is to a user profile;

identifying, based on the user profile, at least one content item of the subset of the plurality of content items, wherein a micro-theme associated with the at least one identified content item comprises the micro-theme of the particular content item;

increasing a relevance value associated with the at least one identified content item of the subset, based on the user interface interaction;

ranking the plurality of content items in an order based at least in part on the respective relevance value associated with each content item of the subset of content items; and

generating for display the plurality of content items based on the ranking.

44. The one or more non-transitory computer-readable media of claim 43, wherein ranking the plurality of content items is further based at least in part on criteria other than the respective relevance values of the subset of the plurality of content items.

45. The one or more non-transitory computer-readable media of claim 43, wherein ranking the plurality of content items is further based at least in part on at least one of date, day, or time.

46. The one or more non-transitory computer-readable media of claim **43**, wherein ranking the plurality of content items is further performed based at least in part on user interface interactions with other user profiles.

47. The one or more non-transitory computer-readable media of claim **43**, wherein ranking the plurality of content items is further performed based at least in part on historical data indicating content items the user profile previously interacted with via an application.

48. The one or more non-transitory computer-readable media of claim **43**, wherein ranking the plurality of content items is further performed based at least in part on a stochastic signature comprising statistical values associated with a plurality of micro-themes, and wherein at least one micro-theme of the plurality of micro-themes is associated with a statistical value indicating its relevance to the user profile.

49. The one or more non-transitory computer-readable media of claim **43**, wherein a trained model is used to:
identify the micro-theme of the particular content item associated with the user interface interaction;
perform the identifying of the at least one content item of the subset of the plurality of content items; and
perform the ranking of the plurality of content items.

50. The one or more non-transitory computer-readable media of claim **43**, wherein the order comprises the plurality

of content items sorted based at least in part on a descending order of relevance, wherein a content item having a higher relevance value has a greater probability of being positioned at a top portion of a display or other prominent portion of the display.

51. The one or more non-transitory computer-readable media of claim **43**, wherein:

the user interface interaction is further associated with a first context comprising at least one of a geographic location, client device identifier, date, day, or time associated with the user interface interaction, and the increasing further comprises:

comparing the first context with a second context associated with at least one content item of the plurality of content items;

determining, based on the comparing, that the second context associated with the at least one content item comprises an attribute of the first context; and

increasing the relevance value associated with the at least one content item.

52. The one or more non-transitory computer-readable media of claim **43**, wherein at least one content item of the plurality of content items comprises at least one of an audio item, a video item, a web content item, an advertisement, or any combination thereof.

* * * * *