(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
20 October 2016 (20.10.2016)

WIPO | PCT

(10) International Publication Number
WO 2016/168703 A1

(51) International Patent Classification:
G06F 17/30 (2006.01) *H04N 7/173* (2011.01)
G06Q 30/00 (2012.01) *H04N 21/25* (2011.01)

(21) International Application Number:
PCT/US2016/027898

(22) International Filing Date:
15 April 2016 (15.04.2016)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/148,668 16 April 2015 (16.04.2015) US

(71) Applicant: **THE REGENTS OF THE UNIVERSITY OF CALIFORNIA** [US/US]; 1111 Franklin Street, 5th Floor, Oakland, CA 94607-5200 (US).

(72) Inventors: **ZHU, Kevin**; 1111 Franklin Street, 5th Floor, Oakland, CA 94607-5200 (US). **CHEN, Wei**; 1111 Franklin Street, 5th Floor, Oakland, CA 94607-5200 (US).

(74) Agents: **ZHANG, Li (Lily)** et al.; Mintz Levin Cohn Ferris Glovsky And Popeo, P.C., 3580 Carmel Mountain Road, Suite 300, San Diego, CA 92130-6768 (US).

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

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

Published:

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

(54) Title: PREFERENCE LEARNING WITH CLUSTERING

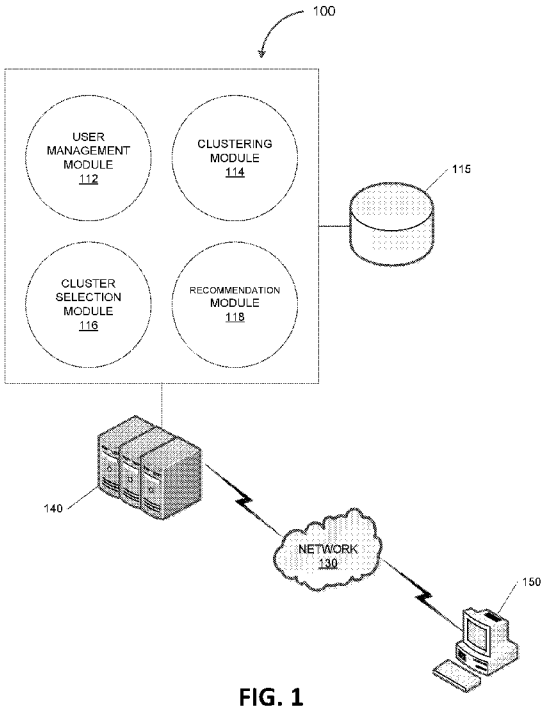


FIG. 1

(57) Abstract: In one aspect, there is provided a method for preference learning. The method may include: determining whether existing data associated with a user is available, wherein the existing data includes preference weights assigned to user clusters; when existing data associated with the user is unavailable: selecting, based on a ranking of the user clusters, a user cluster; providing, based on a selected user cluster, recommendations to the user; and assigning, based on a response of the user to the recommendations, a preference weight to each of the selected user cluster and/or another user cluster; and when existing data associated with the user is available, determining whether to exploit the existing data or to explore new preferences for the user, wherein exploiting comprises providing the recommendations based on the existing data, and wherein exploring comprises providing the recommendations by selecting a user cluster not assigned a preference weight.

WO 2016/168703 A1

PREFERENCE LEARNING WITH CLUSTERING

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to U.S. Provisional Patent Application No. 62/148,668 filed April 16, 2015, entitled "LEARNING HETEROGENEOUS PREFERENCES WITH CLUSTERING ALGORITHMS," the contents of which is hereby incorporated by reference in its entirety.

STATEMENT OF GOVERNMENT SPONSORED SUPPORT

[0002] Certain aspects of the present disclosure were developed with U.S. Government Support under Grant No. NSF IIS-1360566 and NSF IIS-1360568 awarded by the National Science Foundation. The U.S. Government has certain rights in the subject matter of the present disclosure.

TECHNICAL FIELD

[0003] The subject matter disclosed herein relates to machine learning and more specifically to preference learning.

BACKGROUND

[0004] Recommendation systems seek to identify content and/or products that are consistent with individual user preferences. Oftentimes, insights from a recommendation system may be applied towards creating a personalized browsing experience, which can increase a website's traffic and/or sales. For example, a recommendation system may suggest material (e.g., news articles, music, videos) that best interests users of an online content provider.

Similarly, a web retailer may deploy a recommendation system that targets its users with products that are more likely to be purchased by each user.

[0005] Conventional recommendation systems rely primarily on *a priori* knowledge of individual user preferences in order to generate recommendations. That is, recommendations from conventional recommendation systems are generally guided by a user's history of past interactions (e.g., browsing, searches, and/or purchases) with a website. As such, conventional recommendation systems often perform poorly in the absence of existing data such as when encountering new users to a website. Moreover, even when individual user preferences are known, conventional recommendation systems tend to single-mindedly exploit this existing data while foregoing opportunities to discover more through exploration of individual users' potential preferences. Consequently, the performance of conventional recommendation systems may further stagnate over time.

SUMMARY

[0006] Methods, systems, and apparatus, including computer program products, are provided for preference learning. In some example embodiments, there is provided a method that includes: determining whether existing data associated with a user is available, wherein the existing data includes at least one preference weight assigned to at least one of a plurality of user clusters; when existing data associated with the user is not available: selecting, based at least in part on a ranking associated with the plurality of user clusters, a user cluster from the plurality of user clusters; providing, based at least in part on a selected user cluster, one or more recommendations to the user; and assigning, based at least in part on a response of the user to the one or more recommendations, a preference weight to the selected user cluster and/or another user cluster; and when existing data associated with the user is available, determining whether to exploit the existing data or to explore new preferences for the user, wherein exploiting the

existing data comprises providing the one or more recommendations based at least in part on the existing data, and wherein exploring new preferences for the user comprises providing the one or more recommendations at least by selecting a user cluster that is not assigned a preference weight.

[0007] In some variations, one or more of the features disclosed herein including the following features can optionally be included in any feasible combination. The determining may be based on relative proportions such that existing data is exploited during a proportion of occasions and new preferences for the user are explored during a remaining proportion of occasions. The relative proportions may be determined at least by applying one or more of a ε-greedy, ε-first, ε-decreasing, adaptive ε-greedy based on value differences, and contextual ε-greedy strategy. The response of the user may comprise one or more of mouse clicks and purchases of items comprising the one or more recommendations. The preference weight may comprise a vector, and wherein each value in the vector corresponds to a type of response to a recommendation.

[0008] The method may further include selecting, based at least in part on a ranking associated with each of a plurality of user clusters not assigned a preference weight, the user cluster that is not assigned a preference weight. The ranking associated with each of the plurality of user clusters may correspond to one or more of a frequency of interaction with a user cluster, when a user cluster was last selected, and whether a user cluster is subject to a sale and/or promotion.

[0009] The method may further include generating the plurality of user clusters by at least segmenting a plurality of users based at least in part on user attributes and user preferences. The generating of the plurality of user clusters may be performed based on data associated with

other users. Each of the plurality of user clusters may be associated with an item, a category of items, and/or a source of items, and wherein users in each user cluster share a preference for the item, the category or items, and/or the source of items. The method may further include: identifying at least one attribute associated with the user; and assigning a preference weight to a third user cluster based at least in part on a match between the attribute associated with the user and an attribute that is shared by other users in the third user cluster. At least one pair of user clusters comprising the plurality of user clusters may be associated with a proximity weight. The proximity weight may indicate a likelihood that a user having a preference for an item, category of items, and/or source of items that is associated with one user cluster will also have a preference for a different item, category of items, and/or source of items that is associated with another user cluster. The selected user cluster and the other user cluster may be associated with a proximity weight. A preference weight may be assigned to each of the selected user cluster and the other user cluster when the preference weight exceeds a threshold.

DESCRIPTION OF THE DRAWINGS

[00010] The accompanying drawings, which are incorporated in and constitute a part of this specification, show certain aspects of the subject matter disclosed herein and, together with the description, help explain some of the principles associated with the disclosed implementations. In the drawings,

[00011] FIG. 1 depicts a system diagram illustrating a recommendation system, in accordance with some example embodiments;

[00012] FIG. 2 depicts a group of user clusters, in accordance with some example embodiments; and

[00013] FIG. 3 depicts a flowchart illustrating a process for preference learning, in accordance with some example embodiments.

DETAILED DESCRIPTION

[00014] In some example embodiments, a recommendation system may be configured to learn a user's preferences by providing one or more recommendations in the absence of existing data specific to the user. Existing data associated with a user may include the user's attributes and/or preferences. However, such existing data may be unavailable for a new user to a content provider and/or online retailer. As such, the recommendation system may select, based on a ranking associated with the plurality of user clusters, a user cluster from the plurality of user clusters. The recommendation system may provide one or more recommendations based on that user cluster. The recommendation system may assign a preference weight to the user cluster and/or one or more other user clusters based on the user's responses to the recommendations. For example, the preference weight may reflect the user's interactions (e.g., mouse click, purchase) with the products and/or contents indicated by the one or more recommendations.

[00015] In some example embodiments, each individual user cluster may be associated with a preference that is shared by all users included in that user cluster. For example, a user cluster may be associated with an item (e.g., product or content), category (e.g., of products or content), and/or source (e.g., brand, manufacturer, creator of products or content). Users in the user cluster may share a preference for that item, category, and/or source. In the absence of existing data associated with a new user, the recommendation system may select a user cluster based on one or more other criteria including a ranking associated with each user cluster.

[00016] In some example embodiments, the ranking associated with a user cluster may be determined without existing data associated with a specific user. For example, the ranking

associated with a user cluster may correspond to a frequency of previous interactions (e.g. mouse clicks, purchases) by other users with respect to an item, category, and/or source associated with the user cluster. The user cluster ranking may also indicate when the user cluster was last selected. Alternately or additionally, the ranking associated with a user cluster may correspond to whether an item, category, and/or source associated with the user cluster are subject to a sale or promotion.

[00017] In some example embodiments, the recommendation system may generate a plurality of user clusters by segmenting users into groups based on existing data that includes, individual user attributes (e.g., demographics) and preferences (e.g., for an item, category, and/or source). The user clusters may be generated using one or more clustering techniques including, but limited to, connectivity-based clustering, centroid-based clustering, distribution-based clustering, and density-based clustering. Moreover, each individual pair of user clusters may be associated with a proximity weight indicating a relationship between the pair of clusters. For example, the proximity weight may indicate a probability or likelihood that users in one user cluster who have a preference for one item, category, and/or source may also have a preference for a different item, category, and/or source that is associated with another user cluster.

[00018] In some example embodiments, the recommendation system may be configured to continue learning the preferences of a user when the recommendation system has accumulated existing data associated with the user. For instance, the recommendation system may have existing data associated with a user returning to a content provider and/or online retailer. The recommendation system may have existing data that is specific to the user including preference weights that are assigned to one or more user clusters based on the user's past responses to recommendations from the recommendation system. Thus, the recommendation system may

exploit the existing data by selecting, based on the preference weight assigned to one or more user clusters, a user cluster. Alternately or additionally, the recommendation system may exploit the existing data by selecting, based on the proximity weights assigned to at least some pairs of user clusters, a user cluster. The recommendation system may provide one or more recommendations based on the selected user cluster and assign a preference weight to that user cluster based on the user's responses to the recommendations.

[00019] Alternately or additionally, the recommendation system may not exploit the existing data on the user. Instead, the recommendation system may explore new preferences for the user by selecting a user cluster that is not assigned a preference weight. For example, to explore a user's new preferences, the recommendation system may select, from a plurality of user clusters not assigned a preference weight, a user cluster based on the ranking associated with the cluster. The user cluster ranking may reflect when a user cluster was last selected. Thus, the recommendation system may select, based on the ranking, the least (or most) recently selected user cluster. Alternately or additionally, the ranking may correspond to a frequency of interaction with an item, category, and/or source associated with the user cluster. As such, the recommendation system may select a user cluster having the highest (or lowest) frequency of interactions. The user cluster ranking may also correspond to whether an item, category, and/or source associated with the user cluster are subject to a sale and/or promotion. The recommendation system may select, based on the ranking, a user cluster that is (or is not) subject to a sale and/or promotion.

[00020] In some example embodiments, the recommendation system may determine whether to use existing data to provide one or more recommendations to a user. That is, the recommendation system may determine whether to exploit existing data or explore new

preferences for a user. For example, the recommendation system may apply one or more strategies when evaluating the choice between exploitation and exploration. The one or more strategies may include ε-greedy, ε-first, ε-decreasing, adaptive ε-greedy based on value differences, and contextual ε-greedy.

[00021] FIG. 1 depicts a system diagram illustrating a recommendation system 100, in accordance with some example embodiments. Referring to FIG 1, in some example embodiments, the recommendation system 100 may be realized in digital electronic circuitry, integrated circuitry, specially designed application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs) computer hardware, firmware, software, and/or combinations thereof.

[00022] In some example embodiments, the recommendation system 100 may include one or more processors that implement a plurality of modules including a user management module 112, a clustering module 114, a cluster selection module 116, and a recommendation module 116. The recommendation system 100 may include additional and/or different modules without departing from the scope of the present disclosure

[00023] The user management module 112 may be configured to determine the availability of existing data associated with a user 150. For example, the user management module 112 may determine whether existing data on a user is available to the recommendation system 100 when the user starts a new session (e.g., by landing on a website). In some example embodiments, existing data may not be available for a new user to a content provider and/or online retailer. By contrast, at least some existing data (e.g., user attributes and/or preferences) may be available for a user who is returning to a website. For example, the existing data may be available from cookies (e.g., provided by the user), site logs, and/or a

[00024] The clustering module 114 may be configured to generate a plurality of user clusters by segmenting users into groups based on existing data. Each individual user cluster may be associated with a preference that is shared by all users included in that user cluster. For example, a user cluster may be associated with an item (e.g., product or content), category (e.g., of products or content), and/or source (e.g., brand, manufacturer, creator of products or content) that is preferred by every user in that user cluster.

[00025] In some example embodiments, the clustering module 114 may segment, based on existing data on individual user attributes (e.g., demographics) and preferences (e.g., for an item, category, and/or source), users into groups. For example, the clustering module 114 may perform the segmentation of users offline (e.g., separately from, and/or prior to, the provision of any recommendations by the recommendation system 100). In some example embodiments, the clustering module 114 may generate the plurality of user clusters by applying a clustering technique including, for example, connectivity-based clustering, centroid-based clustering, distribution-based clustering, and density-based clustering.

[00026] In some example embodiments, a user cluster may be associated with a preference weight that is specific to a user (e.g., the user 150). The clustering module 114 may assign a preference weight to a user cluster based on the user's response to recommendations made based on that user cluster. The value of the preference weight associated with a user cluster may indicate a likelihood that a user may interact (e.g., click on, purchase) with recommendations made based on that user cluster. For example, a user cluster associated with a specific item (e.g., a smartphone) may be assigned a preference weight having one (e.g., relatively higher) value if a user made a purchase of the specific item. Alternately, the user cluster may be assigned a preference weight having a different (e.g., relatively lower) value if the

user only clicked on the item but did not make a purchase or if the user did not click on the item. In addition, the value of the preference weight associated with the user cluster may be cumulative in order to reflect the user's responses to the item on multiple occasions (e.g., clicks during earlier visits and a purchase during a subsequent visit).

[00027] In some example embodiments, the clustering module 114 may also be configured to assign a proximity weight at least some pairs of user clusters. In some example embodiments, the proximity weight may indicate a relationship between different pairs of user clusters. For example, the value of the proximity weight may correspond to a probability or likelihood that users in one user cluster who have a preference for one item, category, and/or source may also have a preference for a different item, category, and/or source associated with another user cluster. According to some example embodiments, when assigning a preference weight to one user cluster, the clustering module 114 may also assign a preference weight to one or more other user clusters based on proximity weights. For example, when assigning a preference weight to one user cluster, the clustering module 114 may further assign a preference weight to another user cluster if the proximity weight associated with the two user clusters exceeds a threshold.

[00028] In some example embodiments, the clustering module 114 may store data corresponding to the plurality of user clusters, preference weights, and proximity weights in a data store 115 coupled with the recommendation system 100.

[00029] The cluster selection module 116 may be configured to select a user cluster in accordance with the availability of existing data (e.g., as determined by the user management module 112). As such, in some example embodiments, the cluster selection module 116 may be configured to select a user cluster in the absence of existing data associated with a user (e.g., new

user). When existing data associated with a user is unavailable, the cluster selection module 116 may select, based on the ranking associated with individual user clusters, a user cluster. For instance, the cluster selection module 116 may select a user cluster having a highest (or lowest) ranking. The user cluster ranking may be determined without existing data associated with a specific user. Thus, user clustering ranking may be available even a user is new to a website. For example, the user cluster ranking may be determined based on a frequency of previous interactions by other users with respect to an item, category, and/or source associated with the user cluster. Alternately or additionally, the ranking associated with a user cluster may be determined based on when the user cluster was last selected and/or whether the user cluster is subject to a sale or promotion.

[00030] Alternately or additionally, when existing data associated with the user is available, the cluster selection module 116 may be further configured to determine whether to use the existing data to provide one or more recommendations to a user. That is, the cluster selection module 116 may determine whether to exploit the existing data or to explore new preferences for the user. In some example embodiments, the cluster selection module 116 may make the determination based on relative proportions such that the cluster selection module 116 may choose to exploit existing data during a proportion of occasions and explore new preferences during a remaining proportion of occasions. The cluster selection module 116 may determine the relative proportions by applying one or more strategies including ε-greedy, ε-first, ε-decreasing, adaptive ε-greedy based on value differences, and contextual ε-greedy. Accordingly, the cluster selection module 116 may choose to explore new preferences during a ε proportion (e.g., 0.1) of occasions and choose to exploit existing data during a (1- ε) proportion (e.g., 0.9) of occasions.

[00031] In some example embodiments, when the cluster selection module 116 determines to use the existing data to provide one or more recommendations to a user, the cluster selection module 116 may be configured to select a user cluster based on the preference weights that are assigned to user clusters based on the user's responses to recommendations made based on these user clusters. For example, to exploit existing data associated with a user, the cluster selection module 116 may select a user cluster having a highest preference weight.

[00032] In some example embodiments, when the cluster selection module 116 determines to not use the existing data and explore new preferences instead, the cluster selection module 116 may be configured to select a user cluster that is not assigned a preference weight. The cluster selection module 116 may be configured to select, based on proximity weights indicating a relationship to one or more user clusters having a preference weight, a user cluster without a preference weight. The cluster selection module 116 may also select, based on the user cluster ranking, a user cluster without a preference weight. For example, the user cluster ranking may correspond to when the user cluster was last selected. Alternately or additionally, the user cluster ranking may correspond to a frequency of interactions with the user cluster (e.g., clicks on and/or purchases of a corresponding item, category, and/or source). The user cluster ranking may also reflect whether the user cluster is subject to a sale or promotion.

[00033] The recommendation module 116 may be configured to provide, based on a user cluster (e.g., selected by the cluster selection module 116), one or more recommendations. For example, the cluster selection module 116 may select a user cluster that is associated with a specific item (e.g., product or content), category (e.g., of products or content), and/or source (e.g., brand, manufacturer, creator of products or content). As such, the recommendation module

116 may recommend the specific item, items belonging to the category, and/or items originating from the source.

[00034] According to some example embodiments, the recommendation system 110 may be communicatively coupled to a web server 140. The web server 140 may be associated with an entity including, for example, a content provider or online retailer. The web server 140 may be configured to deliver information (e.g., webpages) associated with the entity via a wired and/or wireless network 130 to one or more users including the user 150. In some example embodiments, the information provided by the web server 140 may include recommendations from the recommendation system 100.

[00035] FIG. 2 depicts a group 200 of user clusters, in accordance with some example embodiments. Referring to FIGS. 1-2, the group 200 of user clusters may be generated by the recommendation system 100 (e.g., the clustering module 114) and stored at the data store 115.

[00036] As shown in FIG. 2, the group 200 may include a plurality of user clusters including a first user cluster 210, a second user cluster 220, a third user cluster 230, and a fourth user cluster 240. A group of user clusters (e.g., the group 200 generated by the recommendation system 100) may include a different number of user clusters without departing from the scope of the present disclosure.

[00037] Each of the plurality of user clusters in the group 200 may correspond to a category of items (e.g., products or content). For example, as shown in FIG. 2, the first user cluster 210 may be associated with smartphones, the second user cluster 220 may be associated with tablet personal computers, the third user cluster 230 may be associated with cameras, and the fourth user cluster 240 may be associated with children's apparel. Alternately or additionally, each of the plurality of user clusters in the group 200 may correspond to specific

items (e.g., of products or content) and/or sources (e.g., brands, manufacturers, or creators of products or content) without departing from the scope of the present disclosure.

[00038] The recommendation system 100 (e.g., the clustering module 114) may assign a preference weight to at least some user clusters in the group 200. In some example embodiments, the preference weight assigned to a user cluster may be a value indicative of a user's response to a recommendation made based on the user cluster. For example, the first user cluster 210 may be assigned a preference weight x corresponding to a user's overall response to a recommendation of one or more smartphones. The value of x may indicate a degree of preference and/or a likelihood that the user may interact with (e.g., click on, purchase) recommendations made based on the first cluster 210 (e.g., smartphones).

[00039] Alternately or additionally, in some example embodiments, the preference weight assigned to a user cluster may be a vector representing different types of responses to a recommendation made based on the user cluster. For example, the second user cluster 220 may be assigned with a preference weight vector $[y_1, y_2, \ldots y_n]$. Each value in the preference weight vector $[y_1, y_2, \ldots y_n]$ may correspond to a type of interaction with recommendations made based on the second cluster 220. For example, the value y_1 may indicate a likelihood that the user may click on a recommendation of a tablet personal computer while the value y_2 may indicate a likelihood that the user may make a purchase as a result of a recommendation of a tablet personal computer.

[00040] Some user clusters in the group 200 may not have been used when making a recommendation to a user. For example, the recommendation system 100 may not have selected the third cluster 230 and the fourth user cluster 240, and made any recommendations based on these user clusters. According to some example embodiments, user clusters that have not been

used to make recommendations to a user are not assigned a preference weight. As shown in FIG. 2, the third cluster 230 and the fourth user cluster 240 are not associated with a preference weight since the recommendation system 100 has not made any recommendations of cameras (e.g., based on the third cluster 230) and of children's apparel (e.g., based on the fourth user cluster 240). In some example embodiments, the recommendation system 100 may select a user cluster (e.g., the third user cluster 230, the fourth user cluster 240) that is not associated with a preference weight when the recommendation system 100 is exploring new preferences for a user.

[00041] In some example embodiments, one or more pairs of user clusters may be associated with a proximity weight. The proximity weight may indicate a relationship between the pair of clusters. For example, the first user cluster 210 and the second user cluster 220 may be associated with a proximity weight α that indicates a likelihood that a user who is interested in smartphones (e.g., associated with the first user cluster 210) is also interested in tablet personal computers (e.g., associated with the second user cluster 220), and vice versa. Similarly, the first user cluster 210 and the third user cluster 230 may be associated with a proximity weight β that indicates a likelihood that a user who is interested in smartphones (e.g., associated with the first user cluster 210) is also interested in cameras (e.g., associated with the third user cluster 230). The first user cluster 210 and the fourth user cluster 240 may be associated with a proximity weight γ indicating a likelihood that a user who is interested in smartphones (e.g., associated with the first user cluster 210) is also interested in children's apparel.

[00042] The value of a proximity weight may reflect a degree of likelihood that a user interested in an item, category, and/or source associated with one user cluster is also interested in an item, category, and/or source associated with another user cluster. Thus, the proximity weight

α may be greater than the proximity weight γ if a user who is interested smartphones is more likely to be also interested in tablet personal computers than in children's apparel.

[00043] According to some example embodiments, when assigning a preference weight to one user cluster, the recommendation system 100 (e.g., the clustering module 114) may further assign, based on one or more proximity weights, a preference weight to one or more other user clusters. For instance, the recommendation system 100 may assign the preference weight x (e.g., of the first user cluster 210) based on a user's response to recommendations of smartphones. In addition, the recommendation system 100 may also assign the preference weights for the second user cluster 220 and the third user cluster 230 as a user who demonstrated a interest in smartphones should also be interested in tablet personal computers and cameras as indicated by the proximity weights α and β.

[00044] In some example embodiments, when the recommendation system 100 assigns a preference weight to one user cluster, the recommendation system 100 may also assign a preference weight to another user cluster if the proximity weight associated with both user clusters exceeds a threshold. For example, when the recommendation system 100 assigns the preference weight x to the first user cluster 210, the recommendation system 100 may also assign the preference weight vector $[y_1, y_2, \ldots y_n]$ to the second user cluster 220 if the proximity weight α exceeds a threshold. By contrast, the recommendation system 100 may not assign a preference weight to the third user cluster 230 if the proximity weight β does not exceed the threshold.

[00045] In some example embodiments, when making a recommendation to a user, the recommendation system 100 may select a user cluster based on preference weights and/or proximity weights. For example, the recommendation system 100 may determine to use the existing data associated with a user to provide one or more recommendations. As such, the

recommendation system 100 may select the first user cluster 210 or the second user cluster 220 based on the respective preference weights of each user cluster. In addition to preference weights, the recommendation system 100 may select the first user cluster 210 or the second user cluster 220 based on additional criteria including a ranking associated with each user cluster. The user cluster ranking may correspond to a frequency of interaction (e.g., clicks, purchases) with each user cluster, when a user cluster was last selected, and/or whether a user cluster is subject to a sale or promotion.

[00046] When the recommendation system 100 determines to explore new preferences for the user, the recommendation system may select a user cluster without a preference weight. For example, the recommendation system 100 may select the third user cluster 230 or the fourth user cluster 240 as neither user clusters have been used to provide recommendations to the user and have not been assigned a preference weight indicating the user's response to these recommendations. In some example embodiments, the recommendation system 100 may select a user cluster without a preference weight based on proximity weights relative to user clusters having preference weights. For example, the recommendation system 100 may select the third user cluster 230 based on the proximity weight β with respect to the first user cluster 210 and not the fourth user cluster 240, which is not assigned any proximity weight with respect to a user cluster having a preference weight.

[00047] Alternately or additionally, when exploring new preferences for a user, the recommendation system 100 may select a user cluster without a preference weight based on a ranking of the user cluster. The user cluster ranking may correspond to when a user cluster was last selected. For instance, the recommendation system 100 may select the third user cluster 230 if the third user cluster 230 was selected more (or less) recently than the fourth user cluster. The

user cluster ranking may also reflect a frequency of interaction with the user cluster (e.g., clicks on and/or purchases of a corresponding item, category, and/or source). For example, the recommendation system may select the third user cluster 230 if more (or less) users have purchased and/or clicked on cameras (e.g., associated with the third user cluster 230) than children's apparel (e.g., associated with the fourth user cluster 240). Alternately or addition, the user cluster ranking may correspond to whether the user cluster is subject to a sale or promotion. In one example scenario, the recommendation system 100 may select the third user cluster 230 if an online retailer is running a promotion on cameras.

[00048] In some example embodiments, the recommendation system 100 (e.g., the clustering module 114) may generate the group 200 of user clusters using one or more clustering techniques including finite mixture based cluster, connectivity-based clustering (e.g., personal rank based clustering), centroid-based clustering, distribution-based clustering, and density-based clustering. According to some example embodiments, the group 200 of user clusters may be generated based on existing data including individual user attributes (e.g., demographics) and preferences (e.g., for an item, category, and/or source). Thus, users belonging to the first user cluster 210 may share a common preference for smartphones and one or more common user attributes (e.g., age, gender, income). As such, in some example embodiments, in addition to and instead of a user's responses to recommendations made based on a user cluster, the recommendation system 100 (e.g., the clustering module 114) may assign a preference weight to a user cluster based on a user's attributes. For example, in addition to and instead a user's response to recommendations of smartphones, the recommendation system 100 may assign the preference weight x based on a match between the user's attributes and common user attributes of users in the first user group 210.

[00049] FIG. 3 depicts a flowchart illustrating a process 300 for preference learning, in accordance with some example embodiments. Referring to FIGS. 1-3, the process may be performed by the recommendation system 100.

[00050] The recommendation system 100 may detect a start of a session with a user (302). For example, the recommendation system 100 may detect a start of a session with a user when the user lands on a webpage (e.g. homepage) of an entity (e.g., content provider, online retailer). The session may be a browsing session during which the user interacts with the entity's website including clicking on product and/or content links and making purchases (e.g. of content and/or products).

[00051] The recommendation system 100 may determine whether existing data associated with the user is available (303). For example, the recommendation system 100 may determine whether one or more user clusters have been assigned preference weights associated with the user. In some example embodiments, the recommendation system 100 may determine that existing data associated with the user is not available (303-N). For example, the user may be new to a website, in which case the recommendation system 100 may not have assigned a preference weight to any user clusters based on the user's responses to past recommendations. As such, the recommendation system 100 may select, based at least in part on a ranking of a plurality of user clusters, a user cluster from the plurality of user clusters (304).

[00052] According to some example embodiments, the ranking associated with a user cluster may be determined without existing data (e.g., preference weights) associated with the user. For example, the ranking associated with a user cluster may correspond to a frequency of previous interactions (e.g., mouse clicks, purchases) by other users with respect to an item, category, and/or source associated with the user cluster. A higher (or lower) frequency of

interaction may correspond to a higher (or lower) ranking for a user cluster. The user cluster ranking may also reflect when the user cluster was last selected by the recommendation system 100. A more (or less) recently selected user cluster may have a higher (or lower) ranking than a less (or more) recently selected user cluster. Alternately or additionally, the user cluster ranking may be determined based on whether an item, category, and/or source associated with the user cluster are subject to a sale or promotion. Being subject to a sale or promotion may increase (or decrease) the user cluster ranking.

[00053] The recommendation system 100 may provide, based at least in part on a selected user cluster, one or more recommendations to the user (306). For example, if the recommendation system 100 selects the first user cluster 210, the recommendation system 100 may make one or more recommendations of different smartphones. The recommendation system 100 may make recommendations at any level specificity (e.g., specific products/contents or products/contents from a certain category) without departing from the scope of the present disclosure.

[00054] The recommendation system 100 may assign, based at least in part on a response of the user to the one or more recommendations, a preference weight to the selected user cluster and/or another user cluster (308). For example, based on the user's responses to one or more recommendations of smartphones, the recommendation system 100 may assign the preference weight x to the first user cluster 210. Alternately or additionally, if the user responds to the recommendations by interacting (e.g., clicking on and/or purchasing) a tablet personal computer, the recommendation system 100 may assign the preference weight vector $[y_1, y_2, \ldots y_n]$ to the second user cluster 220.

[00055] In some example embodiments, the preference weight of a user cluster may be a single value (e.g., integer) representative of a user's overall response to recommendations made based on the user cluster. For example, the first user cluster 210 is associated with the preference weight x, which may be assigned without differentiating between the types of responses (e.g., click, purchase) to recommendations (e.g., of smartphones) made based on the first user cluster 210. Alternately or additionally, the preference weight of a user cluster may be a vector representative of different types of user responses to recommendations made based on the user cluster. For example, the second user cluster 220 may be associated with the preference weight vector $[y_1, y_2, \ldots y_n]$. The preference weight vector $[y_1, y_2, \ldots y_n]$ may be assigned depending on the type of responses to the recommendations (e.g., of tablet personal computers) made based on the second user cluster 220.

[00056] In some example embodiments, assigning a preference weight to one user cluster may include assigning a preference weight to another user cluster based on the proximity weight associated with the pair of user clusters. For example, when assigning the preference weight x of the first user cluster 210, the recommendation system 100 may also assign a preference weight to each of the second user cluster 220 and the third user cluster 230 based on the proximity weights α and β. According to some example embodiments, the recommendation system 100 may assign a preference weight to one user cluster when assigning a preference weight to another cluster if the proximity weight associated with the user clusters exceeds a threshold.

[00057] The recommendation system 100 may determine whether the user has ended the session (309). If the recommendation system 100 determines that the user has ended the session (309-Y), the recommendation system 100 may end the process 300 (310). For example, the user

may end the session by leaving the website. In some example embodiments, upon the end of a session with a user, the recommendation system 100 may be configured to store data collected during the user's session (e.g., preference weights) such that the recommendation system 100 may have existing data associated with the user during subsequent sessions with the user. Alternately or additionally, if the recommendation system 100 determines that the user has not ended the session (309-N), the recommendation system 100 may determine whether to use existing data associated with the user to provide one or more recommendations to the user (311).

[00058] Alternately or additionally, in response to detecting the start of a session with a user at operation 302, the recommendation system 100 may determine that existing data associated with the user is available (303-Y). For instance, the recommendation system 100 may identify existing data associated with the user if the user is a user returning to the website. As such, the recommendation system 100 may determine whether to use the existing data associated with the user to provide one or more recommendations to the user (311). That is, the recommendations system 100 may determine whether to exploit the existing data or explore new preferences for the user.

[00059] In some example embodiments, the recommendation system 100 may apply one or more strategies when determining whether to use the existing data to provide one or more recommendations to the user. The one or more strategies may include ε-greedy, ε-first, ε-decreasing, adaptive ε-greedy based on value differences, and contextual ε-greedy. The recommendation system 100 may apply the one or more strategies to determine the relative proportions of occasions during which the recommendation system 100 may choose to exploit existing data versus explore new preferences for the user. For instance, the recommendation

system 100 may choose to explore new preferences during a ε proportion (e.g., 0.1) of occasions and choose to exploit existing data during a (1- ε) proportion (e.g., 0.9) of occasions.

[00060] If the recommendation system 100 determines to use existing data associated with the user to provide one or more recommendations to the user (311-Y), the recommendation system 100 may select, based at least in part on the preference weight associated with one or more user clusters, a user cluster (312). For example, the recommendation system 100 may select the first user cluster 210 if the first user cluster 210 has a higher preference weight than the second user cluster 220. In addition to preference weights, in some example embodiments, the recommendation system 100 may select a user cluster based on a ranking associated with the user cluster. For example, the user cluster ranking may reflect a frequency of previous interactions by other users with respect to an item, category, and/or source associated with the user cluster. The user cluster ranking may also correspond to when the user cluster was last selected and/or whether the user cluster is subject to a sale or promotion.

[00061] In some example embodiments, the recommendation system 100 may provide, based at least in part on a selected user cluster, one or more recommendations to the user 306). For example, the recommendation system 100 may recommend a specific item (e.g., a smartphone) associated with a user cluster. Alternately or additionally, the recommendation system 100 may also recommend items in a category of items (e.g., electronics) and/or originates from a source (e.g., brand, manufacturer, creator) associated with a user cluster. The recommendation system 100 may further assign, based at least in part on a response of the user to the one or more recommendations, a preference weight to the selected user cluster and/or another user cluster (308). For example, if the recommendation system 100 selects the first user cluster 210, the recommendation system 100 may assign a preference weight x to the first user

cluster 210 based on the user's responses to recommendations of smartphones. The recommendation system 100 may also assign preference weights to the second user cluster 220 and/or the third user cluster 230 based on the user's responses to those user clusters and/or based on the relationship between these user clusters and the first user cluster 210 as indicated by the proximity weights α and β.

[00062] Alternately or additionally, if the recommendation system 100 determines to not use the existing data associated with the user to provide one or more recommendations to the user (311-N), the recommendation system 100 may explore new preferences for the user. Accordingly, the recommendation system 100 may select a user cluster that is not associated with a preference weight (314). When the recommendation system 100 determines to explore new preferences for the user, the recommendation system 100 may select user clusters that have not been previously used in making recommendations and thus have not been assigned any preference weights. For example, the recommendation system 100 may select the third user cluster 230 or the fourth user cluster 240, which have not been assigned with any preference weights.

[00063] In some example embodiments, the selection of a user cluster without preference weights may be made based on proximity weights with respect to user clusters that have been assigned preference weights. For example, the recommendation system 100 may select the third user cluster 230 and not the fourth user cluster 240 based on the proximity weight β relating the third user cluster 230 to the first user cluster 210, which has been assigned the preference weight x. Alternately or additionally, the recommendation system 100 may select a user cluster without a preference weight based on a user cluster ranking. For example, the user cluster ranking may correspond to a frequency of previous interactions by other users with

respect to an item, category, and/or source associated with the user cluster. The user cluster ranking may also correspond to when the user cluster was last selected and/or whether the user cluster is subject to a sale or promotion.

[00064] The recommendation system 100 may provide, based at least in part on a selected user cluster, one or more recommendations to the user (306). For example, the recommendation system 100 may recommend a specific item (e.g., a smartphone) associated with a user cluster. Alternately or additionally, the recommendation system 100 may also recommend items in a category of items (e.g., electronics) and/or originates from a source (e.g., brand, manufacturer, creator) associated with a user cluster. The recommendation system 100 may further assign, based at least in part on a response of the user to the one or more recommendations, a preference weight to the selected user cluster and/or another user cluster (308). For instance, the recommendation system 100 may assign a preference weight to a user cluster based on a user's interactions (e.g., clicks, purchases) with a recommended item, category of items, and/or source for items. The recommendation system 100 may further assign preference weights to one or more other user clusters based on proximity weights.

[00065] The recommendation system 100 may determine whether the user has ended the session (309). If the recommendation system 100 determines that the user has ended the session (309-Y), the recommendation system 100 may end the process 300 (310). Alternately, if the recommendation system 100 determines that the user has not ended the session (309-N), process 300 may continue at operation 311.

[00066] The process 300 may include additional and/or different operations than shown without departing from the scope of the present disclosure. One or more operations of the

process 300 may be omitted and/or repeated without departing from the scope of the present disclosure.

[00067] One or more aspects or features of the subject matter described herein can be realized in digital electronic circuitry, integrated circuitry, specially designed application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs) computer hardware, firmware, software, and/or combinations thereof. These various aspects or features can include implementation in one or more computer programs that are executable and/or interpretable on a programmable system including at least one programmable processor, which can be special or general purpose, coupled to receive data and instructions from, and to transmit data and instructions to, a storage system, at least one input device, and at least one output device. The programmable system or computing system may include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[00068] These computer programs, which can also be referred to as programs, software, software applications, applications, components, or code, include machine instructions for a programmable processor, and can be implemented in a high-level procedural and/or object-oriented programming language, and/or in assembly/machine language. As used herein, the term "machine-readable medium" refers to any computer program product, apparatus and/or device, such as for example magnetic discs, optical disks, memory, and Programmable Logic Devices (PLDs), used to provide machine instructions and/or data to a programmable processor, including a machine-readable medium that receives machine instructions as a machine-readable signal. The term "machine-readable signal" refers to any signal used to provide machine

instructions and/or data to a programmable processor. The machine-readable medium can store such machine instructions non-transitorily, such as for example as would a non-transient solid-state memory or a magnetic hard drive or any equivalent storage medium. The machine-readable medium can alternatively or additionally store such machine instructions in a transient manner, such as for example, as would a processor cache or other random access memory associated with one or more physical processor cores.

[00069] To provide for interaction with a user, one or more aspects or features of the subject matter described herein can be implemented on a computer having a display device, such as for example a cathode ray tube (CRT) or a liquid crystal display (LCD) or a light emitting diode (LED) monitor for displaying information to the user and a keyboard and a pointing device, such as for example a mouse or a trackball, by which the user may provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well. For example, feedback provided to the user can be any form of sensory feedback, such as for example visual feedback, auditory feedback, or tactile feedback; and input from the user may be received in any form, including acoustic, speech, or tactile input. Other possible input devices include touch screens or other touch-sensitive devices such as single or multi-point resistive or capacitive track pads, voice recognition hardware and software, optical scanners, optical pointers, digital image capture devices and associated interpretation software, and the like.

[00070] The subject matter described herein can be embodied in systems, apparatus, methods, and/or articles depending on the desired configuration. The implementations set forth in the foregoing description do not represent all implementations consistent with the subject matter described herein. Instead, they are merely some examples consistent with aspects related to the described subject matter. Although a few variations have been described in detail above,

other modifications or additions are possible. In particular, further features and/or variations can be provided in addition to those set forth herein. For example, the implementations described above can be directed to various combinations and subcombinations of the disclosed features and/or combinations and subcombinations of several further features disclosed above. In addition, the logic flows depicted in the accompanying figures and/or described herein do not necessarily require the particular order shown, or sequential order, to achieve desirable results. Other implementations may be within the scope of the following claims.

WHAT IS CLAIMED IS:

1. A method for preference learning, comprising:

determining whether existing data associated with a user is available, wherein the existing data includes at least one preference weight assigned to at least one of a plurality of user clusters;

when existing data associated with the user is not available:

selecting, based at least in part on a ranking associated with the plurality of user clusters, a user cluster from the plurality of user clusters;

providing, based at least in part on a selected user cluster, one or more recommendations to the user; and

assigning, based at least in part on a response of the user to the one or more recommendations, a preference weight to the selected user cluster and/or another user cluster; and

when existing data associated with the user is available, determining whether to exploit the existing data or to explore new preferences for the user, wherein exploiting the existing data comprises providing the one or more recommendations based at least in part on the existing data, and wherein exploring new preferences for the user comprises providing the one or more recommendations at least by selecting a user cluster that is not assigned a preference weight.

2. The method as recited in claim 1, wherein the determining is based on relative proportions such that existing data is exploited during a proportion of occasions and new preferences for the user are explored during a remaining proportion of occasions.

3. The method as recited in claim 2, wherein the relative proportions are determined at least by applying one or more of a ε-greedy, ε-first, ε-decreasing, adaptive ε-greedy based on value differences, and contextual ε-greedy strategy.

4. The method as recited in claim 1, further comprising selecting, based at least in part on a ranking associated with each of a plurality of user clusters not assigned a preference weight, the user cluster that is not assigned a preference weight.

5. The method as recited in claim 1, wherein the ranking associated with each of the plurality of user clusters corresponds to one or more of a frequency of interaction with a user cluster, when a user cluster was last selected, and whether a user cluster is subject to a sale and/or promotion.

6. The method as recited in claim 1, further comprising generating the plurality of user clusters by at least segmenting a plurality of users based at least in part on user attributes and user preferences.

7. The method as recited in claim 6, wherein the generating of the plurality of user clusters is performed based on data associated with other users.

8. The method as recited in claim 6, wherein each of the plurality of user clusters is associated with an item, a category of items, and/or a source of items, and wherein users in each user cluster share a preference for the item, the category or items, and/or the source of items.

9. The method as recited in claim 6, further comprising:

identifying at least one attribute associated with the user; and

assigning a preference weight to a third user cluster based at least in part on a match between the attribute associated with the user and an attribute that is shared by other users in the third user cluster.

10. The method as recited in claim 1, wherein at least one pair of user clusters comprising the plurality of user clusters is associated with a proximity weight.

11. The method as recited in claim 10, wherein the proximity weight indicates a likelihood that a user having a preference for an item, category of items, and/or source of items that is associated with one user cluster will also have a preference for a different item, category of items, and/or source of items that is associated with another user cluster.

12. The method as recited in claim 10, wherein the selected user cluster and the other user cluster are associated with a proximity weight.

13. The method as recited in claim 12, wherein a preference weight is assigned to each of the selected user cluster and the other user cluster when the preference weight exceeds a threshold.

14. The method as recited in claim 1, wherein the response of the user comprises one or more of mouse clicks and purchases of items comprising the one or more recommendations.

15. The method as recited in claim 1, wherein the preference weight comprises a vector, and wherein each value in the vector corresponds to a type of response to a recommendation.

16. A system for preference learning, comprising:

at least one processor; and

at least one memory including program code which when executed by the at least one memory provides operations comprising:

determining whether existing data associated with a user is available, wherein the existing data includes at least one preference weight assigned to at least one of a plurality of user clusters;

when existing data associated with the user is not available:

selecting, based at least in part on a ranking associated with the plurality of user clusters, a user cluster from the plurality of user clusters;

providing, based at least in part on a selected user cluster, one or more recommendations to the user; and

assigning, based at least in part on a response of the user to the one or more recommendations, a preference weight to the selected user cluster and/or another user cluster; and

when existing data associated with the user is available, determining whether to exploit the existing data or to explore new preferences for the user, wherein exploiting the existing data comprises providing the one or more recommendations based at least in part on the existing data, and wherein exploring new preferences for the user comprises providing the one or more recommendations at least by selecting a user cluster that is not assigned a preference weight.

17. The system as recited in claim 16, wherein the determining is based on relative proportions such that existing data is exploited during a proportion of occasions and new preferences for the user are explored during a remaining proportion of occasions.

18. The system as recited in claim 17, wherein the relative proportions are determined at least by applying one or more of a ε-greedy, ε-first, ε-decreasing, adaptive ε-greedy based on value differences, and contextual ε-greedy strategy.

19. The system as recited in claim 16, further comprising selecting, based at least in part on a ranking associated with each of a plurality of user clusters not assigned a preference weight, the user cluster that is not assigned a preference weight.

20. The system as recited in claim 16, wherein the ranking associated with each of the plurality of user clusters corresponds to one or more of a frequency of interaction with a user cluster, when a user cluster was last selected, and whether a user cluster is subject to a sale and/or promotion.

21. The system as recited in claim 16, further comprising generating the plurality of user clusters by at least segmenting a plurality of users based at least in part on user attributes and user preferences.

22. The system as recited in claim 21, wherein the generating of the plurality of user clusters is performed based on data associated with other users.

23. The system as recited in claim 21, wherein each of the plurality of user clusters is associated with an item, a category of items, and/or a source of items, and wherein users in each user cluster share a preference for the item, the category or items, and/or the source of items.

24. The system as recited in claim 21, further comprising:

identifying at least one attribute associated with the user; and

assigning a preference weight to a third user cluster based at least in part on a match between the attribute associated with the user and an attribute that is shared by other users in the third user cluster.

25. The system as recited in claim 16, wherein at least one pair of user clusters comprising the plurality of user clusters is associated with a proximity weight.

26. The system as recited in claim 25, wherein the proximity weight indicates a likelihood that a user having a preference for an item, category of items, and/or source of items that is associated with one user cluster will also have a preference for a different item, category of items, and/or source of items that is associated with another user cluster.

27. The system as recited in claim 25, wherein the selected user cluster and the other user cluster are associated with a proximity weight.

28. The system as recited in claim 27, wherein a preference weight is assigned to each of the selected user cluster and the other user cluster when the preference weight exceeds a threshold.

29. The system as recited in claim 16, wherein the response of the user comprises one or more of mouse clicks and purchases of items comprising the one or more recommendations.

30. The system as recited in claim 16, wherein the preference weight comprises a vector, and wherein each value in the vector corresponds to a type of response to a recommendation.

31. A non-transitory computer-readable storage medium including program code which when executed by at least one processor causes operations comprising:

determining whether existing data associated with a user is available, wherein the existing data includes at least one preference weight assigned to at least one of a plurality of user clusters;

when existing data associated with the user is not available:

selecting, based at least in part on a ranking associated with the plurality of user clusters, a user cluster from the plurality of user clusters;

providing, based at least in part on a selected user cluster, one or more recommendations to the user; and

assigning, based at least in part on a response of the user to the one or more recommendations, a preference weight to the selected user cluster and/or another user cluster; and

when existing data associated with the user is available, determining whether to exploit the existing data or to explore new preferences for the user, wherein exploiting the existing data comprises providing the one or more recommendations based at least in part on the existing data, and wherein exploring new preferences for the user comprises providing the one or more recommendations at least by selecting a user cluster that is not assigned a preference weight.

32. An apparatus for preference learning, comprising:

means for determining whether existing data associated with a user is available, wherein the existing data includes at least one preference weight assigned to at least one of a plurality of user clusters;

when existing data associated with the user is not available:

means for selecting, based at least in part on a ranking associated with the plurality of user clusters, a user cluster from the plurality of user clusters;

means for providing, based at least in part on a selected user cluster, one or more recommendations to the user; and

means for assigning, based at least in part on a response of the user to the one or more recommendations, a preference weight to the selected user cluster and/or another user cluster; and

when existing data associated with the user is available, means for determining whether to exploit the existing data or to explore new preferences for the user, wherein exploiting the existing data comprises providing the one or more recommendations based at least in part on the existing data, and wherein exploring new preferences for the user comprises providing the one or more recommendations at least by selecting a user cluster that is not assigned a preference weight.

33. The apparatus as recited in claim 32, further comprising means for performing the method as recited in any of claims 2-15.

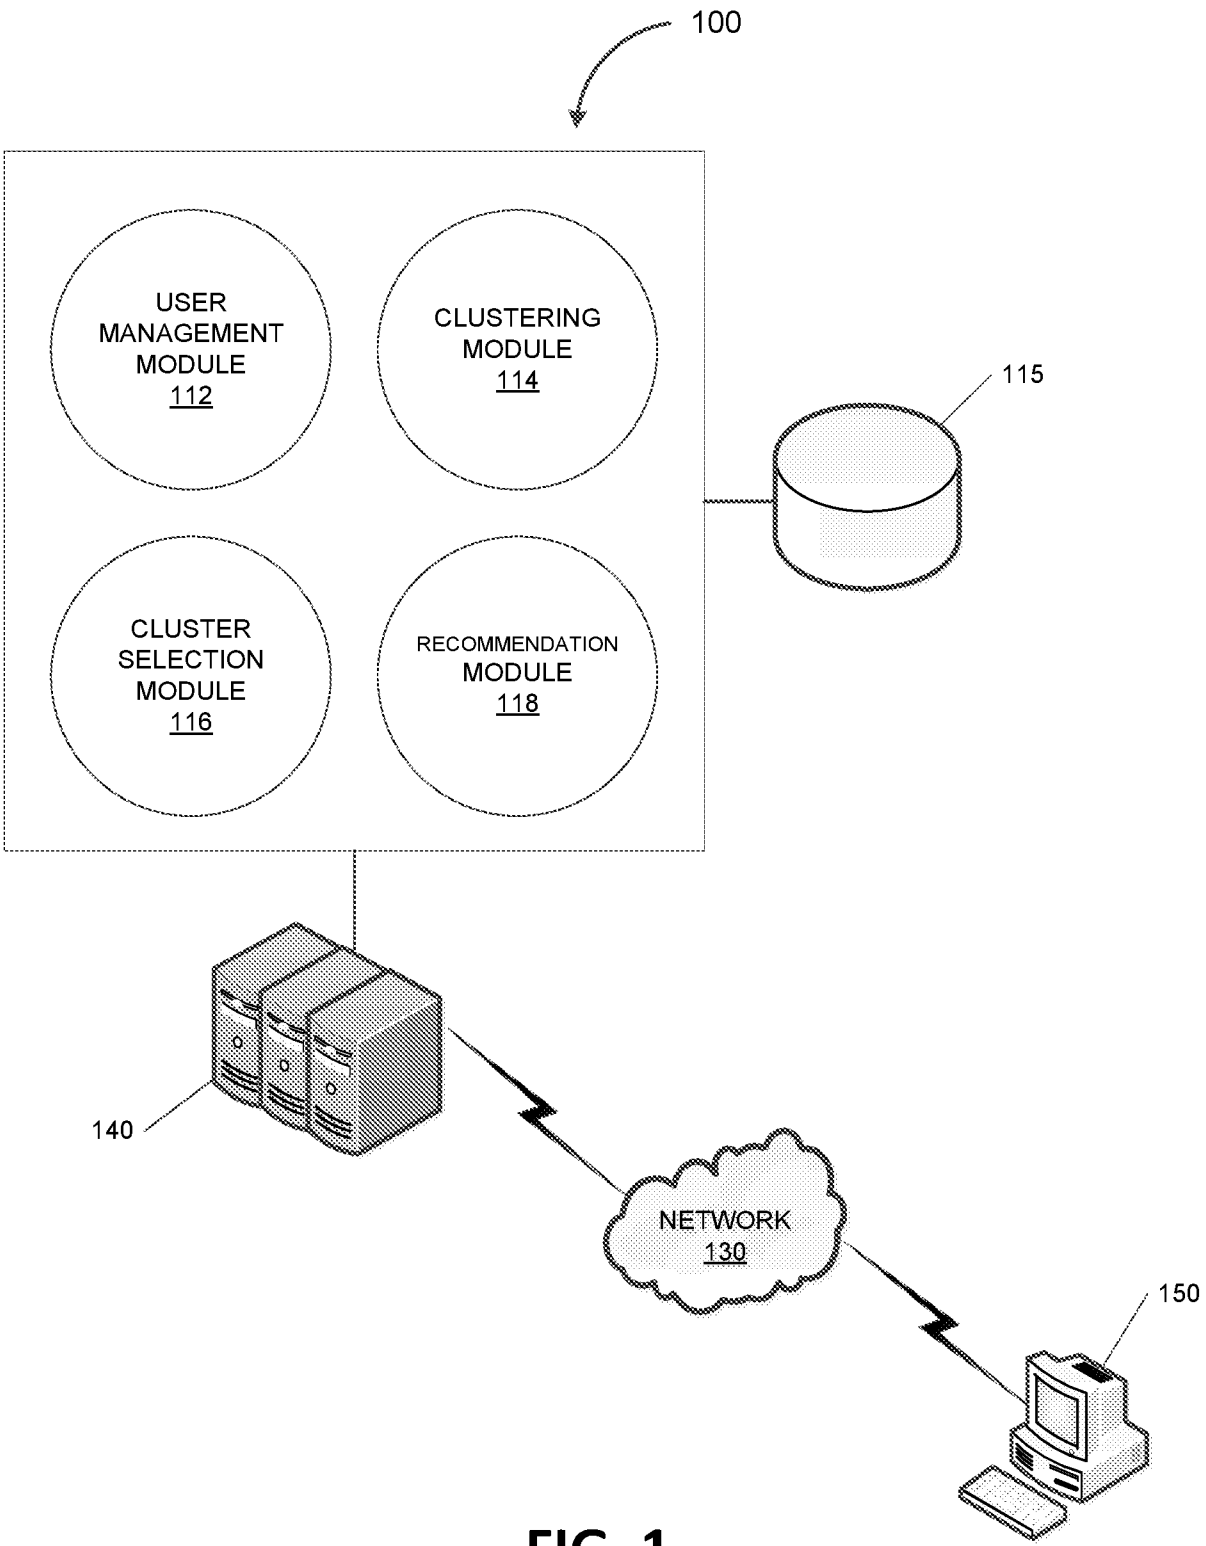

100
USER MANAGEMENT MODULE 112
CLUSTERING MODULE 114
CLUSTER SELECTION MODULE 116
RECOMMENDATION MODULE 118
115
140
NETWORK 130
150

FIG. 1

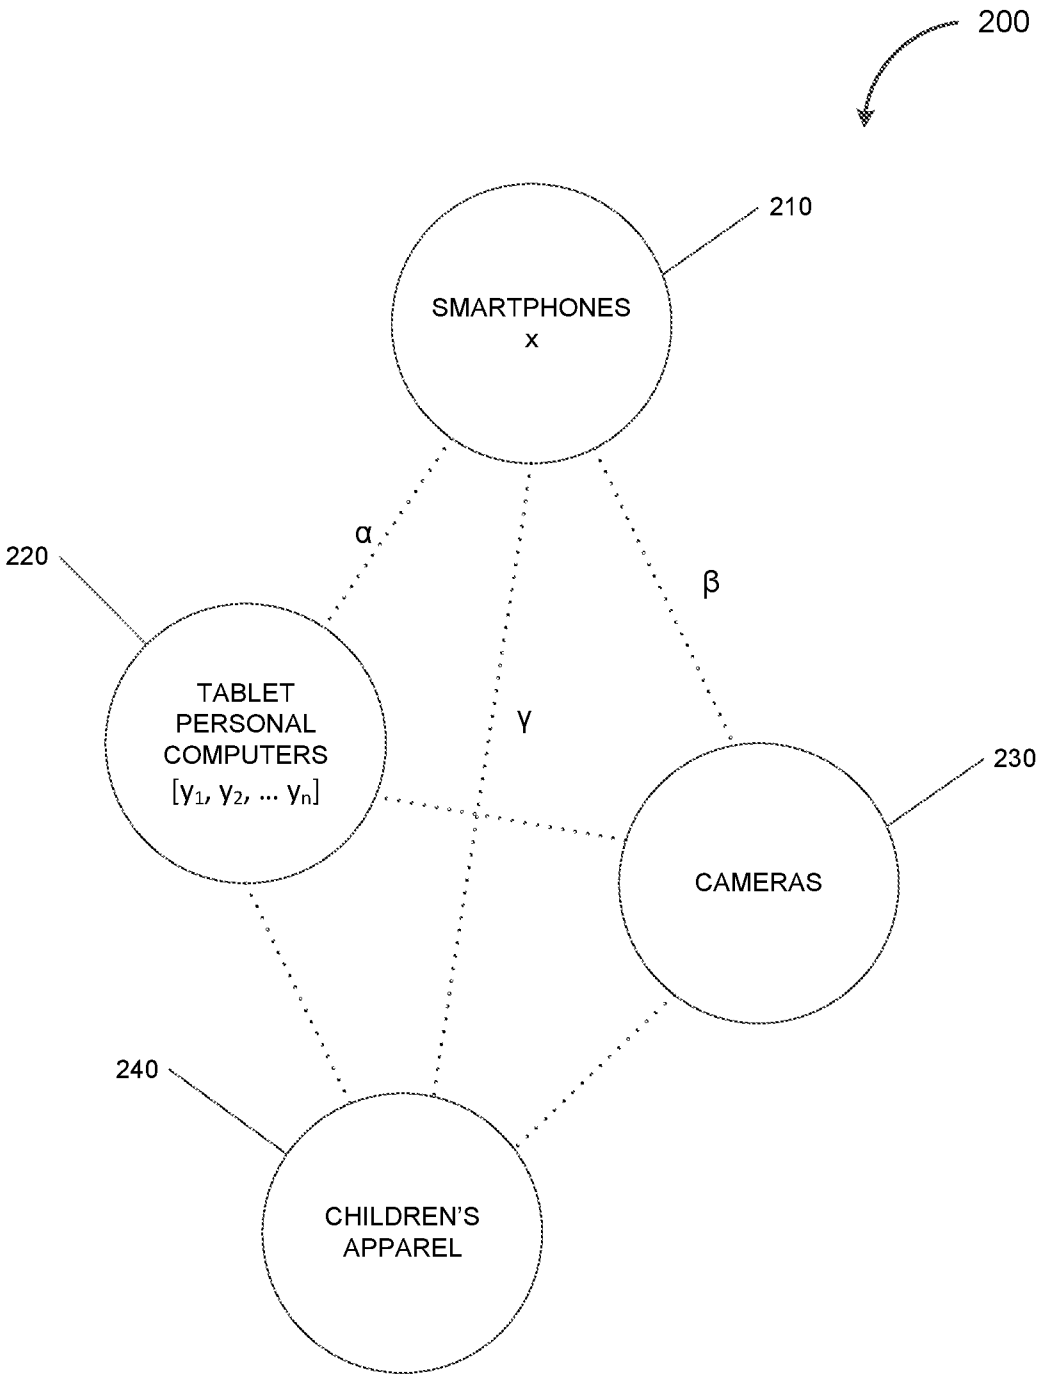

200
210
SMARTPHONES
x
α
β
γ
220
TABLET
PERSONAL
COMPUTERS
[y1, y2, ... yn]
230
CAMERAS
240
CHILDREN'S
APPAREL

FIG. 2

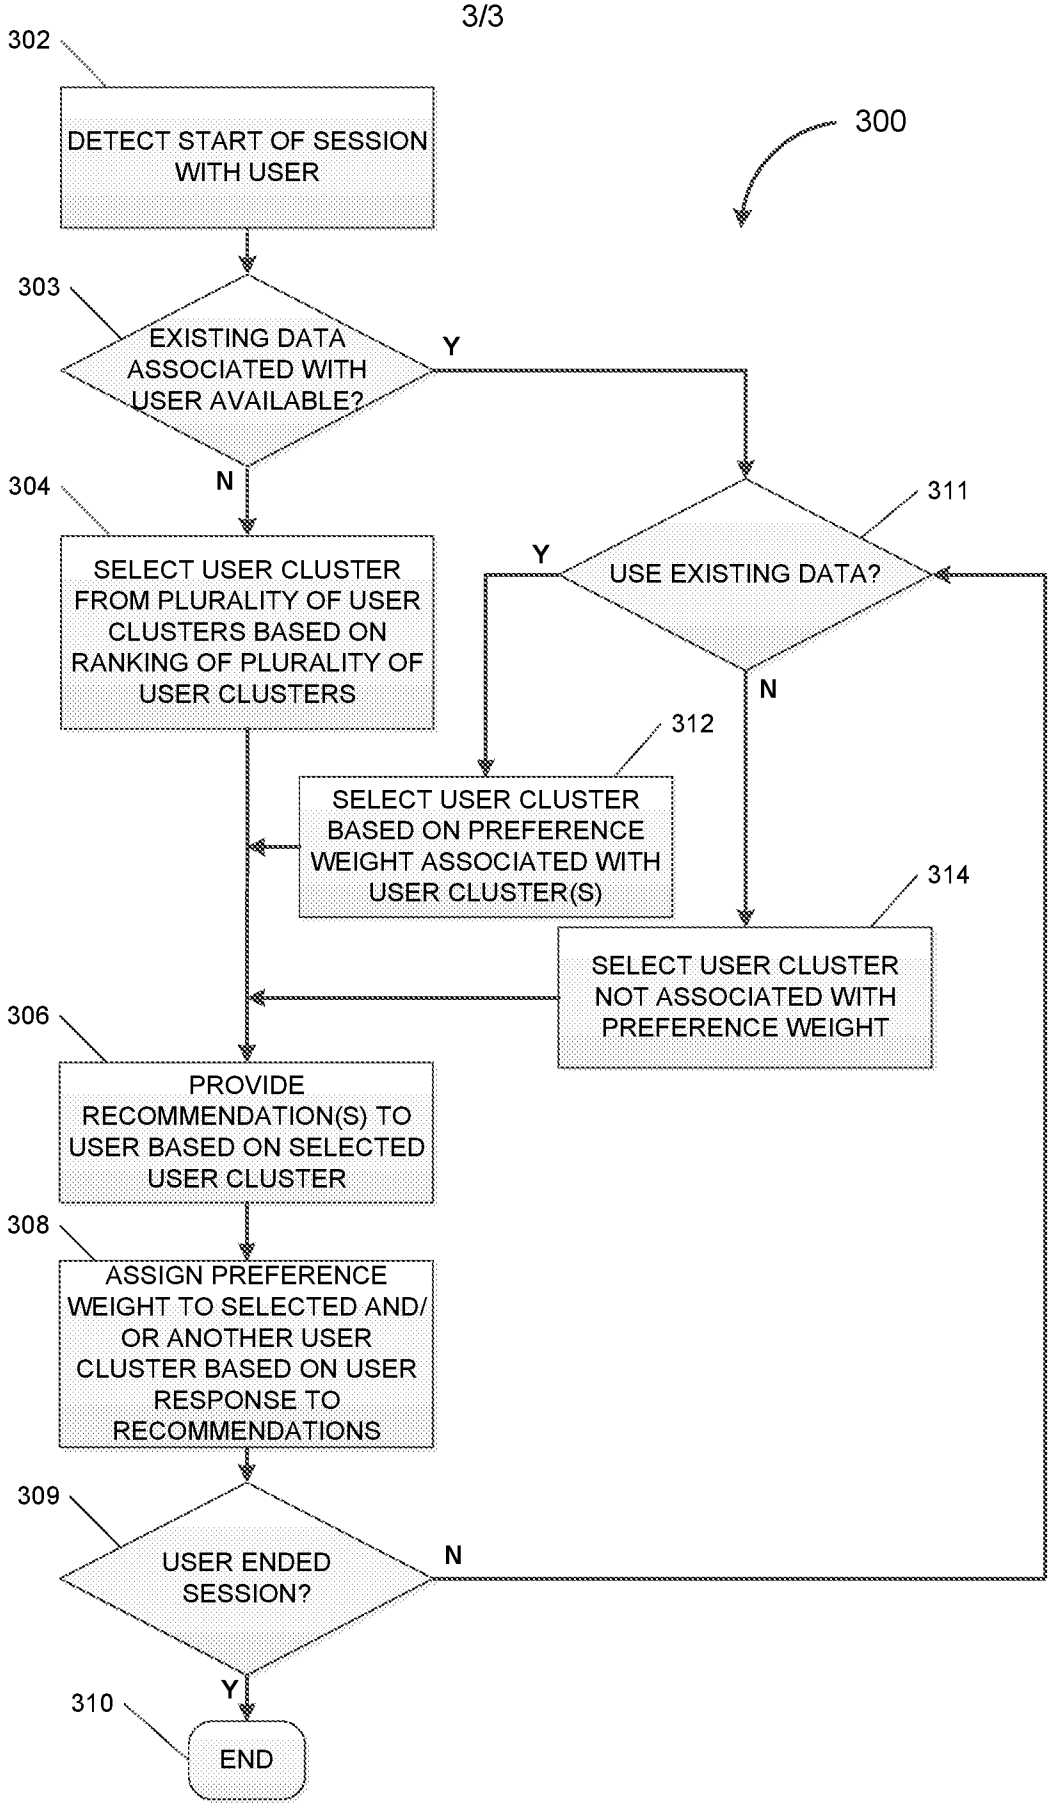

300
302
DETECT START OF SESSION WITH USER
303
EXISTING DATA ASSOCIATED WITH USER AVAILABLE?
Y
N
304
SELECT USER CLUSTER FROM PLURALITY OF USER CLUSTERS BASED ON RANKING OF PLURALITY OF USER CLUSTERS
311
USE EXISTING DATA?
Y
N
312
SELECT USER CLUSTER BASED ON PREFERENCE WEIGHT ASSOCIATED WITH USER CLUSTER(S)
314
SELECT USER CLUSTER NOT ASSOCIATED WITH PREFERENCE WEIGHT
306
PROVIDE RECOMMENDATION(S) TO USER BASED ON SELECTED USER CLUSTER
308
ASSIGN PREFERENCE WEIGHT TO SELECTED AND/OR ANOTHER USER CLUSTER BASED ON USER RESPONSE TO RECOMMENDATIONS
309
USER ENDED SESSION?
N
Y
310
END

FIG. 3

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/027898

A. CLASSIFICATION OF SUBJECT MATTER
IPC(8) - G06F 17/30, G06Q 30/00, H04N 7/173, H04N 21/25 (2016.01)
CPC - G06F 17/30734, G06Q 30/02, H04N 7/173, H04N 21/252, H04N 21/44213, H04N 21/4661 (2016.02)
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC(8) - G06F 17/30, G06Q 30/00, H04N 7/173, H04N 21/25 (2016.01)
CPC - G06F 17/30734, G06Q 30/02, H04N 7/173, H04N 21/252, H04N 21/44213, H04N 21/4661, H04N 21/4666, H04N 21/4668
H04N 21/47202, H04N 21/6582 (2016.02)

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC - 706/12, 707/794, 707/E17.099, 725/9 (keyword delimited)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
Orbit, Google Patents, Google Scholar, Google
Search terms used: preference learning, whether existing data associated with user available,

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	PARK et al. "Pairwise preference regression for cold-start recommendation." In: Proceedings of the third ACM conference on Recommender systems. 23 October 2009 (23.10.2009) Retrieved from <http://www.gatsby.ucl.ac.uk/~chuwei/paper/p21-park.pdf>, entire document	1-33
Y	US 2015/0039539 A1 (TELEFONAKTIEBOLAGET L M ERICSSON (PUBL)) 05 February 2015 (05.02.2015), entire document	1-33
Y	VANCHINATHAN et al. "Explore-exploit in top-n recommender systems via gaussian processes." In: Proceedings of the 8th ACM Conference on Recommender systems. 06 October 2014 (06.10.2014) Retrieved from <https://www.researchgate.net/profile/Hastagiri_Prakash_Vanchinathan/publication/277330674_Explore-Exploit_in_Top-N_Recommender_Systems_via_Gaussian_Processes/links/556c3e2008aeab7772216b8a.pdf>, entire document	2, 3, 17, 18
Y	US 2009/0164516 A1 (SVENDSEN et al) 25 June 2009 (25.06.2009), entire document	10-13, 25-28
A	WO 2010/144766 A1 (HUNCH INC.) 16 December 2010 (16.12.2010), entire document	1-33
A	US 2013/0066921 A1 (MARK et al) 14 March 2013 (14.03.2013), entire document	1-33
A	US 2012/0036523 A1 (WEINTRAUB et al) 09 February 2012 (09.02.2012), entire document	1-33
A	US 2014/0108307 A1 (WIPRO LIMITED) 17 April 2014 (17.04.2014), entire document	1-33

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
15 June 2016	15 JUL 2016
Name and mailing address of the ISA/ Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, VA 22313-1450 Facsimile No. 571-273-8300	Authorized officer Blaine R. Copenheaver PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/210 (second sheet) (January 2015)