

(51) International Patent Classification:
G06F 17/30 (2006.01)(21) International Application Number:
PCT/US2012/046536(22) International Filing Date:
12 July 2012 (12.07.2012)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/506,752 12 July 2011 (12.07.2011) US

(72) Inventors; and

(71) Applicants : PINCKNEY, Thomas [US/US]; 55 W. 25th Street, #27D, New York, New York 10010 (US). DIXON, Christopher [US/US]; 79 E. 2nd Street, New York, New York 10003 (US). GATTIS, Matthew Ryan [US/US]; 45 E. 25th Street, New York, New York 10010 (US).

(74) Agents: MADDEN, Robert B., Reg. No. 57,521 et al.; Schwegman, Lundberg & Woessner, P.A., P.O. Box 2938, Minneapolis, Minnesota 55402 (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, 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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, 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, 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, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: RECOMMENDATIONS IN A COMPUTING ADVICE FACILITY




Fig. 50

(57) Abstract: According to various embodiments, a ratings matrix including matrix values is generated, each row of the ratings matrix identifying one of a plurality of users, each column of the ratings matrix identifying one of a plurality of items, and each of the matrix values corresponding to a known affinity rating describing a degree of affinity associated with one of the users and one of the items. The ratings matrix may include a missing entry representing an unknown affinity rating. According to various embodiments, a revised ratings matrix is generated by factoring the ratings matrix into a user matrix and an item matrix, the revised ratings matrix being the product of the user matrix and the item matrix and including at least one entry representing a predicted affinity rating in place of the missing entry.

WO 2013/010024 A1



Published:

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

RECOMMENDATIONS IN A COMPUTING ADVICE FACILITY

CROSS REFERENCE TO RELATED APPLICATION

[0001] This PCT application claims the benefit of the filing date of U.S. Provisional Patent Application Serial No. 61/506,752, filed July 12, 2011 entitled, "RECOMMENDATIONS IN A COMPUTING ADVICE FACILITY," the entire content of which is incorporated herein by reference.

BACKGROUND

Field:

[0002] The present application relates generally to providing recommendations regarding different items to different users.

Description of the Related Art:

[0003] Online searching for topical advice represents a significant use of computer resources such as provided through the Internet. Computer users may currently employ a variety of search tools to search for advice on specific topics, but to do so may require expertise in the use of search engines, and may produce voluminous search results that take time to sift through, interpret, and compare. People may be accustomed to asking other people for advice in spoken natural language, and therefore it may be useful to have a computer-based advice tool that mimics more closely how people interact with each other. In addition, advice on topics may change in time, and any static database of advice may fall quickly out of date. Therefore, a need exists for improved topical advice search capabilities adapted for use with natural language, and that provides for continuous content refinement.

BRIEF DESCRIPTION OF THE FIGURES

[0004] The invention and the following detailed description of certain embodiments thereof may be understood by reference to the following figures:

[0005] Fig. 1 depicts a list of topics in the system from which users may get decisions.

- [0006] Fig. 2 depicts an example question that the system may ask a user.
- [0007] Fig. 3 depicts an example picture question that the system may ask a user.
- [0008] Fig. 4 depicts an example of the type of information the system may show the user when making a particular decision.
- [0009] Fig. 5 depicts an example of top lists for cameras.
- [0010] Fig. 6 depicts a second example of a top list for cameras.
- [0011] Fig. 7 depicts an embodiment of a user home page.
- [0012] Fig. 8 and 8A depict an embodiment of a user's remembered answers.
- [0013] Fig. 9 depicts choices that a user may contribute expertise.
- [0014] Fig. 10 depicts an example of a user question.
- [0015] Fig. 11 and 11A depict an embodiment of an answer format.
- [0016] Fig. 12 depicts an example list of all decisions in a topic.
- [0017] Fig. 13 depicts an embodiment process flow for the present invention.
- [0018] Fig. 14 depicts an embodiment process flow for the present invention.
- [0019] Fig. 15 depicts an embodiment of a block diagram for the present invention.
- [0020] Fig. 16 depicts an embodiment contributor / expert interface home page.
- [0021] Fig. 17 depicts an embodiment objective question to a user looking for help in a decision.
- [0022] Fig. 18 depicts an embodiment of a decision result showing a particular recommended decision.
- [0023] Fig. 19 depicts an embodiment interface for users to set associations between attributes and decision results.
- [0024] Fig. 20 depicts an embodiment illustrating how a user may edit a decision result.
- [0025] Fig. 21 depicts an embodiment showing prior revisions to content and changes between two prior revisions.

- [0026] Fig. 22 depicts an embodiment showing a question being edited by a user.
- [0027] Fig. 23 depicts an embodiment showing the revision history for attributes.
- [0028] Fig. 24 depicts an embodiment of a workshop interface where newly added areas of advice may be displayed.
- [0029] Fig. 25 depicts an embodiment where the system is asking the user a subjective question in order to learn the preferences of the user.
- [0030] Fig. 26 depicts an embodiment showing an activity feed of recent activity by contributors.
- [0031] Fig. 27 depicts an embodiment showing results based on multiple dimensions.
- [0032] Fig. 28 depicts an embodiment showing multiple question and answer results in response to a user's unstructured input.
- [0033] Fig. 29 depicts an embodiment showing an example question to the user, asking for their preference.
- [0034] Fig. 30 depicts a similarity profile of the news personality Glenn Beck as determined in embodiments of the present invention.
- [0035] Fig. 31 depicts a similarity profile of the personality Martha Steward as determined in embodiments of the present invention.
- [0036] Fig. 32 depicts an embodiment for using a third-party API to help learn about a user.
- [0037] Fig. 33 depicts an embodiment for using a third-party API to help learn about a user and to target responses returned to the user from a user query.
- [0038] Fig. 34 depicts an embodiment for determining an unknown user's preferences through the use of Internet social interactive graphical representations.
- [0039] Fig. 35 depicts an embodiment for the improvement of user taste and preference profiling.
- [0040] Fig. 36 depicts an embodiment of a web-based advice facility interfacing with a graph construct.
- [0041] Fig. 37 depicts an embodiment of an interestingness recommendation process block diagram.

[0042] Fig. 38 depicts an embodiment of a local discovery application visual representation of recommendations to a user.

[0043] Fig. 39 depicts an embodiment of a local discovery application visual representation of linked detail for a recommendation to a user.

[0044] Fig. 40 depicts an embodiment of a local discovery application visual representation of recommendations to a user.

[0045] Fig. 41 depicts an embodiment of a review of a restaurant.

[0046] Fig. 42 depicts an embodiment of an interestingness recommendation process flow diagram.

[0047] Fig. 43 depicts an embodiment of a geographically localized recommendation process flow diagram.

[0048] Fig. 44 is a block diagram of an example system, according to various embodiments.

[0049] Fig. 45 illustrates examples of portions of affinity ratings information, according to various embodiments.

[0050] Fig. 46 illustrates an example of a ratings matrix, according to various embodiments.

[0051] Fig. 47 illustrates an example of a user matrix, according to various embodiments.

[0052] Fig. 48 illustrates an example of an item matrix, according to various embodiments.

[0053] Fig. 49 illustrates an example of a revised ratings matrix, according to various embodiments.

[0054] Fig. 50 is a flowchart illustrating an example method, according to various embodiments.

[0055] Fig. 51 illustrates an example of a ratings matrix, according to various embodiments.

[0056] Fig. 52 illustrates an example of a user matrix, according to various embodiments.

[0057] Fig. 53 illustrates an example of a ratings matrix, according to various embodiments.

[0058] Fig. 54 illustrates an example of an item matrix, according to various embodiments.

[0059] Fig. 55 is a flowchart illustrating an example method, according to various embodiments.

[0060] Fig. 56 is a flowchart illustrating an example method, according to various embodiments.

[0061] Fig. 57 is a flowchart illustrating an example method, according to various embodiments.

[0062] Fig. 58 is a flowchart illustrating an example method, according to various embodiments.

[0063] Fig. 59 illustrates an example of a confidence matrix, according to various embodiments.

[0064] Fig. 60 is a flowchart illustrating an example method, according to various embodiments.

[0065] Fig. 61 is a flowchart illustrating an example method, according to various embodiments.

[0066] Fig. 62 is a flowchart illustrating an example method, according to various embodiments.

[0067] Fig. 63 is a flowchart illustrating an example method, according to various embodiments.

[0068] Fig. 64 is a flowchart illustrating an example method, according to various embodiments.

[0069] Fig. 65 is a diagrammatic representation of machine in the example form of a computer system within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed.

[0070] While the invention has been described in connection with certain preferred embodiments, other embodiments would be understood by one of ordinary skill in the art and are encompassed herein.

[0071] All documents referenced herein are hereby incorporated by reference.

DETAILED DESCRIPTION

[0072] The present invention may ask the user 1314 questions 1320 and based on the user's answers the system may offer a decision, such as a recommendation, a diagnosis, a conclusion, advice, and the like. Internally, the

system may use machine learning to optimize which questions 1320 to ask and what decision 1310 to make at the end of the process. The system may learn through users giving feedback on the ultimate decision, whether the decision 1310 was helpful or not. Helpful solutions may get reinforced and associated with the questions 1320 and answers 1322 that were asked along the way. When a user 1314 says that a decision 1310 was helpful the system may remember which questions 1320 it asked, what the answer 1322 to each question 1320 was, and may associate these questions 1320 and answers 1322 with the ultimate decision. These associations may be the basis of the machine learning that learns over time which question 1320 to ask the next time a user 1314 comes to the system.

[0073] For example a user 1314 may try to get advice picking a bar to visit. The system may ask the question "How old are you?" and get the answer "in my 30s". Ultimately, the system may show the user 1314 the decision "Kelley's Irish Bar". Assume the user 1314 says this decision was helpful. The system will increase the association between the question "How old are you?", the answer "in my 30s" and the decision "Kelley's Irish Bar". The next time a user 1314 comes to the site looking for advice on a bar, the system will be more likely to ask the user 1314 the "How old are you?" question 1320 since in the past this question 1320 was useful in helping the user. If the user 1314 answers the question 1320 in the same way as the prior user 1314 (saying "in my 30s") then the system will increase its belief that the ultimate decision is "Kelley's Irish Pub".

[0074] The system may build a profile of each user's tastes, aesthetic preferences, etc. and learn via feedback which decisions 1310 are liked by which types of people. Alternatively, an expert user may specify which kinds of decisions 1310 are liked by which kinds of people. Learning user's taste profiles may happen through a separate process from the dialog of questions 1320 and answers 1322 asked by the system in a specific topic. For example, a user 1314 may separately tell the system about their taste choices through a different question and answer dialog designed specifically to understand the user's aesthetic preferences.

[0075] A user 1314 may not want to spend the time to teach the system about all of their taste preferences, and so instead the system may learn, or an

expert may specify, which of all the taste questions 1320 are the most important taste questions to ask in the context of the user 1314 making one specific decision 1310. Out of the universe of all questions the system may know about finding out about taste profiles, for instance the system may have learned there are three specific questions 1310 that are best for when the user 1314 is trying to find a sedan under \$25,000. Alternately, there may be a completely different set of three taste questions to ask a user 1314 who is interested in a SUV over \$45,000.

[0076] A user 1314 may also only tell the system about their taste preferences and not about any objective questions. In this case the system may provide a ranking of all the decisions 1310 in an area of advice based purely on taste. So instead of the user 1314 saying they want a \$200 point-and-shoot camera, effectively what the user 1314 would be doing is saying they want a camera that other urban 35 year old men who prefer computers to sports want. Users 1314 may indicate this preference by using a search interface and choosing an area of advice that is explicitly labeled "cameras for urban men in their 30s" instead of the "which camera should I buy" area of advice. Alternatively, users 1314 may indicate their interest in making a decision 1310 about cameras and then opt to not answer any of the questions in the Q&A dialog from the system and thus the system will only have subjective information about the user 1314 to use in recommending cameras to the user 1314. Alternatively, users 1314 may answer questions 1320 in the dialog that are both objective and subjective and the system may then recommend a camera based on the combined objective data about the camera and subjective data about the camera.

[0077] Users may also enter new questions, answers, and ultimate decisions. The system may then try out the new questions 1320 with future users to see if the questions 1320 turn out to be useful in helping those users. For example, a user 1314 of the bar recommendation service may contribute the question "Do you want a loud place or a quiet intimate setting?". The system may decide to ask this question 1320 in a future use of the bar recommendation service and through the process outlined above observe a correlation between the answers of this question 1320 and recommendations that users find useful. On the other hand, a user 1314 may contribute a question 1320 that has no value in

helping users. For example, a user 1314 could contribute the question "Do you have a Canon camera?". The system may try this question 1320 out on future users and fail to notice any correlation between the answers to this question 1320 and bar recommendations that users find helpful. In this case, the question 1320 may get asked less since it's not predictive of whether one recommendation or another recommendation is helpful.

[0078] The system may keep asking questions 1320 until it feels it has a high confidence in a few possible decisions. The system may also stop sooner if it feels like it has already asked too many questions 1320 and risks annoying the user. The system may also ask at least a minimum number of questions 1320 to avoid the user 1314 feeling that the system couldn't possibly have asked enough to make an intelligent decision.

[0079] The system may have a mechanism to tolerate incorrect answers from the user. Incorrect answers may result from the user 1314 not understanding the question, not understanding the answer 1322 or not knowing the answer 1322 to the question. If the bulk of the answers given by the user 1314 support a particular decision, the system may make that decision 1310 even though not all the user's answers support that decision.

[0080] In embodiments, the present invention may provide for at least one of questions 1320 and answers 1322 between the system and the user, decisions to users, and machine learning utilized to improve decisions. The system may provide for an improved way to generate questions 1320 and answers 1322, an improved way to provide decisions to users, an improved way to utilize machine learning to improve questions 1320 and decisions provided by a system, and the like, where any of these capabilities may be separately, or in combination, used as a standalone system or incorporated into a third party system as an improved capability. In embodiments, each of these improved capabilities may utilize some form of machine learning as described herein. For example, the system may provide for an improved way to execute a question 1320 and answer 1322 session with a user 1314 by learning under what circumstances the user 1314 is looking for certain information. For instance, it may be learned by the system that the weather is a condition under which users have a differentiated preference depending on the time of day and the weather conditions. When it's raining during the day, and a user 1314 searches for

movies, the user 1314 may be more likely to be looking for movie tickets and locations where the movie is playing. When it's raining during the night, and the user 1314 searches for movies, the user 1314 may be more likely to be looking for a description of the movie. In another example, the system may provide for an improved way to provide decisions to users, such as learning that users prefer certain formats during the daytime versus during the evening, providing choices verses a single decision 1310 based on age, prefer a greater number of questions 1320 prior to presentation of the decision 1310 based on the user's geographic location, and the like. In another example, the system may provide for an improved way to learn what decision 1310 to choose for a user, such as utilizing greater expert information based on age and education, utilizing popular opinion more when the topic is fashion and the user 1314 is young versus utilizing traditional practice more when the user 1314 is older, asking more questions 1320 about the user's choices in friends when the topic is personal, and the like.

[0081] In embodiments, the present invention may provide for combinations of question 1320 and answer, providing decisions, and learning what decisions to provide, where one of the elements may not be provided by the system, such as when that element is provided by a third party system. For example, a third party search engine web application may want to improve their capabilities for providing sorted lists from a user's search query, and so may want to utilize the present invention's facility for generating questions 1320 and answers 1322 to augment their keyword search and sort algorithms. In this instance, the third party search engine provider may not be interested in the present invention's facility for generating decisions, because their service is in the business of providing sorted lists, not a limited set of decisions. However, the present invention may provide an important new capability to the search engine provider, in that the present invention's ability to constantly improve the questions 1320 and answers 1322 to users may enable the search engine provider to improve their sorting result to users based on the present invention's capabilities.

[0082] In embodiments, the subject of the initial area of advice may be specified through a search interface. For example, a user 1314 searching for "romantic honeymoons in Italy" may get taken to a web page that helps the user 1314 decide where to honeymoon in Italy instead of first asking the user

questions about where they want to vacation, what type of vacation they were looking for etc. Or a user 1314 could search for a specific location in Italy and be directed to a web page on that 1) helps the user 1314 decide if that specific location is a good one for their needs (for example, showing things like "this vacation is good for honeymooners and romantic getaways and bad for family vacations") and 2) offers to start a dialog to help the user 1314 find alternative and potentially better locations in Italy to vacation. Or a user 1314 could be searching for specific products and then enter into a dialog to narrow down which of those products are best for them. In both cases # 1 and #2 the information shown may be based how other users have answered questions in decision making dialogs and then given positive feedback to this decision. So if many people using the "where should I go on vacation" topic answered a question "do you want a romantic vacation" with "yes" and then gave positive feedback to "Italy" the system would display that Italy is a romantic destination to users 1314 coming in via search engines. Alternatively, the users 1314 who added the decision "Italy" or the question "do you want a romantic vacation" into the system could have explicitly indicated that the answer "yes" to the question "do you want a romantic vacation" should be associated with Italy and thus show that Italy is a romantic vacation to users 1314 coming in via search engines.

[0083] In embodiments, the present invention may provide other combinations of some subset of asking questions, making decisions, and learning to make better decisions, such as using the present invention's facilities for making better decisions, but only using input from experts; not providing a question 1320 and answer 1322 session for a particular user, but rather utilize previous user 1314 interactions with the system to provide decisions; asking questions 1320 and answers 1322 to a user 1314 to allow the system to learn in association with future decisions, but providing rewards to the user 1314 rather than decisions; asking question 1320 and answers and making a decision 1310 without any learning, such as simply filtering down results; utilizing the present invention's ability to learn how to make a better decision, but providing that capability to an expert system rather than to users through a question 1320 and answer 1322 interface; and the like. In embodiments, the system may provide for all the elements of a question 1320 and answer 1322 user 1314 session, providing decisions to the user, and learning how to improve decisions.

[0084] In embodiments, a user 1314 entering a question 1320 may optionally specify dependencies and importances for the question. Dependencies may control when the question can be asked. Importances may specify relative importances between different questions 1320 for weighing a user's 1314 answers. If the system has to make trade-offs because no one decision 1310 result matches all of the answers 1322 specified by a user 1314, the system may try to recommend decision results that match high importance questions over lower importance questions. The system may also prioritize asking high importance questions over low importance questions. For example, a user 1314 entering a new question like "Where in the United States do you want to vacation" set a dependency that requires an existing question such as "Where in the world do you want to go" to have been answered with "The United States" before the new question "Where in the United States do you want to vacation" can be asked.

[0085] In embodiments, the present invention may provide for a system with a user 1314 interface through which the user 1314 may interact with the facilities of the system. The system may include several parts, some of which may be the website, the supervisor, and a collection of widgets. Widgets may be collections of code that collect, process, and render a single piece of content on the website. The website may consist of interfaces for end-users, staff members, and registered users to get decisions, edit the decisions, and view reports on system performance. The supervisor may be a container for running widgets so that a widget can perform time-consuming data collection and processing ahead of user 1314 requests to render that content.

[0086] For example, a widget might collect videos about decisions from the internet. The widget, in the supervisor, might crawl the web looking for videos about each decision 1310 and store videos it finds in a database. When the user 1314 comes to the website and gets a particular decision, the website may ask the video widget to render itself and display any videos it has previously found.

[0087] A plurality of instances of the supervisor may be running on multiple computers in order to scale up the widget's processing. Each widget may be running on its own computer. Similarly, many computers may be providing interfaces to the system through web-servers, instant messaging, voice

gateways, email, programmatic APIs, via being embedded in third party websites, and the like.

[0088] In embodiments, attributes may be combinations of a question 1320 and one particular answer 1322 to that question. For example, if a question 1320 was "How old are you?" and the answers to that question 1320 were "under 18", "20 - 30" and "over 30", then an attribute would be "How old are you? Under 18". The system may work by learning the relationship between attributes and decisions. When the system asks a question 1320 and the user 1314 gives an answer 1322 then the system may take that attribute and see which decisions are associated with it.

[0089] In embodiments, the system may understand that some attributes represent continuous values while others represent discrete values. When using continuous attributes, the system may be able to make more intelligent tradeoffs such as understanding that it is frequently acceptable to recommend a product that costs less than the user 1314 asked for but rarely acceptable to offer a product that costs more than the user 1314 asked for.

[0090] In embodiments, the relationships between attributes and decisions may be learned from users, explicitly given to the system or some combination of the two, and the like. For example, a price attribute of "How much do you want to spend? Under \$200" might be explicitly linked to cameras that fall into that price range based on data from experts, ecommerce sites/APIs, etc. The relationship between the attribute "How will you use the camera? On vacations" and possible vacation destinations might be fully learned however.

[0091] When entering new questions 1320, answers 1322, and results the user 1314 may optionally specify the relationships between attributes and decision results. For example, if a user 1314 were to enter the question "how much do you want to spend?" in the "which camera should I buy" topic, the user 1314 may also specify to the system that the answer "under \$200" should be associated with cameras X and Y but not camera Z. Then, if a future user were to use the "which camera should I buy" topic and were to answer the "how much do you want to spend" question with the answer "under \$200" that user 1314 may have a higher chance of being recommended camera X and Y over camera Z.

[0092] After seeking advice from the system and receiving a decision result, a user 1314 may also be given reasons from the system as to why that particular decision result was recommended. This explanation may also allow the user 1314 to change the attributes for the decision result if the user 1314 believes that the decision result was recommended in error by the system.

[0093] In general, the relationships learned may involve training from users, experts, employees, automated data feeds from third parties, or some combination.

[0094] In embodiments, there may be various ways that the system can recommend a solution and select the next question 1320 to ask the user. Possible machine learning systems may be geometric systems like nearest neighbors and support vector machines, probabilistic systems, evolutionary systems like genetic algorithms, decision trees, neural networks, associated with decision trees, Bayesian inference, random forests, boosting, logistic regression, faceted navigation, query refinement, query expansion, singular value decomposition and the like. These systems may be based around learning from complete game plays (e.g., all attributes given by a user 1314 before getting a decision), the answers to individual questions / subsets of game plays, only positive feedbacks, only negative feedbacks or some combination of the two. Additionally, the system may take into account previous interactions the user 1314 had such as remembering previously answered questions, decisions that the user 1314 liked or did not like, which areas of advice the user 1314 previously sought advice in, etc. Additionally, the system may take into account factors that are implicitly provided by the user 1314 such as time of day and date the user 1314 used the system, the user's IP address, client type (e.g., Firefox, IE, cell phone, SMS, and the like), and other such data.

[0095] In embodiments, the present invention may provide for a machine learning system that goes well beyond the capabilities of collaborative filtering, such as through explicitly asking questions 1320 instead of implicitly learning based on a user's behavior, which may be much more powerful since the system is not left trying to infer the user's intent, mood, etc. Also, choosing the questions 1320 to ask the user 1314 based on what they've already answered may allow the present invention to zero in on nuances that would otherwise be missed. The present invention may have the ability to explain decisions, such as providing

decisions beyond simple extrapolations from past behavior such as in, 'other people who bought X, Y and Z also liked product A'. Instead, the present invention may be able to say the user 1314 should 'do A because the user 1314 said they wanted X, liked Y and believed Z'. In addition, the present invention may allow users to contribute new questions 1320 that may be useful, and then automatically learn under which contexts, if any, the question 1320 is helpful. In another area of difference, the present invention's machine learning technology may be able to provide decisions in a great variety of user 1314 interest areas, where collaborative filtering has difficulties being applied to non-product / media applications. For instance, collaborative filtering would not be easily applied to helping a user 1314 make a decision 1310 on a highly personal topic, such as whether they should get a tattoo, or a rare question 1320 such as whether a particular expense can be deducted on the user's tax return. The present invention may be capable of such applications. In embodiments, the present invention may be able to use pre-programmed expert advice inter-mixed with advice learned from a group of users to make decisions to users.

[0096] In embodiments, the system may have a wiki web interface for editing all of the data on the system. The web interface may be used to edit/create/delete questions, answers, attributes, and solutions. Each solution may also have a variety of information associated with it, which may be shown on the decision page when that solution is recommended. For example, when recommending a vacation in Cancun the recommendation page might show videos about Cancun. All of this ancillary data about the solution may also be editable through the wiki.

[0097] In embodiments, the wiki may be used to edit data collected by widgets running in the supervisor. This may allow the widgets to collect data ahead of time and then have a human quality assurance process to review and change the collected data.

[0098] In embodiments, the system may maintain a history of all changes made by either the widgets or humans. For example, one use of this history may be to review the work done by hired contractors doing content quality assurance. Another use of this history may be making sure that the widgets do not undo work done by humans. For example, if the widgets collect a particular video and a human deletes that video because it is inappropriate, then the widget can use

the history to not re-add that video again sometime in the future. Finally, if data is corrupted or incorrectly deleted the history may allow a means of recovery.

[0099] In embodiments, when widgets find new content they may queue tasks to a human workflow for validating and editing that content.

[00100] In embodiments, in order to learn, the system may sometimes make random or semi-random decisions in hopes of recommending something that the system wouldn't have expected to be useful, but which may turn out to be useful. If the system wants to use what it has already learned, then it may not make random choices in which questions 1320 it asks and which decision 1310 it makes. There may be a tradeoff between using what is already known, also referred to as exploitation, and potentially learning something new, also referred to as exploration. Exploitation may lead to a more satisfied user, while exploration may make the system smarter.

[00101] In embodiments, one way to make this trade-off when selecting questions 1320 to ask the user 1314 may be to ask questions 1320 that the system is confident are useful in making a decision 1310 and then picking a few random questions 1320 to ask. Another way to make the trade-off may be to have a fixed budget in every user 1314 interaction where a fixed set of questions 1320 are based on exploitation and the next set are based on exploration.

[00102] In embodiments, decisions may also be explored or exploited. If the system wants to learn, it may show a random decision. Instead of showing a purely random decision, the system may also show a decision 1310 that meets some requirements specified by the user 1314 and is purely exploring within the remaining requirements. For example, instead of picking a random camera to show the user 1314 the system could pick a random camera that meets the user's price requirements. This may result in more efficient training since the system may be less likely to show a decision 1310 that has no chance of meeting the user's needs. Rather than showing a random decision 1310 when exploring, the system may also show both the exploited decision 1310 and an explored solution and get feedback on each separately from the user. Alternatively, the system could inject a limited amount of randomness and pick a decision "like" what the system's best guess is. For example, the system may predict that the user 1314 will like one particular camera but could instead recommend another similar but not identical camera in order to balance making a reasonable decision 1310 and

still learning new information from the user. In embodiments, the system may identify to the user 1314 when it is asking questions 1320 or making decisions through exploration vs. exploitation, or it may not.

[00103] In embodiments, the system may be viewed as surveying users about the various things it is recommending. For example, the system may ask the user 10 questions 1320 about the Canon SD1000 camera. This may provide a rich set of data about each camera allowing the system to start building lists of what kind of user 1314 is likely to like this camera. The system may build a ranked list of decisions for each attribute, such as from most likely to be liked to least likely to be liked, given that attribute. For example, the system may build a list of cameras in order likely to be liked by people who say "How old are you? Over 50". This may be shown by the system as the top 10 cameras for users over 50. Numerous of these top 10 lists may be constructed based on the system's data. These lists may also be combined to form new lists. For example, given the ranked lists of cameras for the attribute "How old are you? Over 50" and another list for the attribute "Why are you buying a camera? Travel", the system may construct a new ranked list of cameras for the "Over 50 year old users who want a travel camera". These combinations of top lists may be pre-generated, generated on-demand by incrementally asking the user 1314 to select new top lists, and the like.

[00104] In embodiments, these "top lists" may be used for a variety of purposes. Some users may not want to answer a series of questions 1320 before receiving a decision. Instead, they may be able to browse through these lists and find a relevant decision. The system may have a large number of top lists, such as thousands or tens of thousands, each of which may have their own web page. In addition, these pages may contain a large amount of content that may be indexed by search engines and bring users to the system's website. Alternatively, users 1314 may use a search interface in the system itself to find the area of advice they want a decision in. Various top lists may be used to short-cut the dialog by implicitly answering some of the questions 1320 in the dialog based on the toplist. For example, there could be an area of advice called "vacations" and a top list called "romantic honeymoon vacations in Italy" that servers as a short cut or gateway into the "vacations" topic with several questions 1320 from the "vacations" dialog already answered: "Where do you want to go? Europe",

"Where in Europe do you want to go? Italy", "Are you traveling on a special occasion? Yes", "What is the special occasion? honeymoon". This may serve as an alternate interface for the user 1314 to seek advice through a traditional search interface without engaging in a question and answer dialog.

[00105] In embodiments, various pages on the site may have self-contained displays of information called widgets. For example, the decision pages may have a widget that shows how other people who liked this question 1320 answered various questions, videos/pictures about the decision, links to other web sites that have information about the decision, personalized pros and cons of this decision 1310 based on how the user 1314 answered questions, lists of other decisions that a similar, lists of other decisions that would have been made had questions 1320 been answered differently, lists of awards/honors for this decision (such as Consume Reports recommended), and the like.

[00106] In embodiments, the system may allow users to navigate through the universe of decisions (e.g., cameras, vacation destinations, etc) along dimensions that are not commonly available. For example, instead of being shown a camera and only letting the user 1314 say "show me more/less expensive cameras" the system may let the user 1314 say "show me cameras that are more liked by young people", "show me a camera that is better for travel and less stylish", and the like. Dimensions like "style", "good for travel", "bad for young people", and the like, may be generated as a side-effect by asking users questions 1320 and then learning what is a good decision 1310 given those answers.

[00107] In embodiments, navigating along alternative dimensions may be used as a starting point for the user 1314, instead of the user 1314 selecting an area to seek advice in and then engaging in a dialog. The user 1314 may start interacting with the system by using a search interface or an external search engine to search for a specific decision result, such as a product name or travel destination. The system would then show the user information about that specific decision result and allow the user 1314 to navigate to other decision results, engage in a dialog to refine what the user 1314 is looking for, or show the user 1314 information that the system has learned (through machine learning, expert advice or some combination) about this specific decision result. For example, a user 1314 may use a search interface to navigate to a web page showing

information on a Canon SD1100 camera. The system may show other cameras that people looking for a Canon SD1100 also like, allow the user 1314 to find similar cameras along non-traditional feature dimensions such as a camera that is better for taking pictures of sporting events, as well as show what the system knows about the Canon SD1100 such as "great for travel", "not good for people learning photography", "Available for under \$200", "Preferred by people who are consider themselves design conscious", and the like.

[00108] In embodiments, another possible interface may be to show users a list of decisions and display a simple explanation for why each decision 1310 is being made. For example, when recommending cameras the system may show three cameras and say that one is "cheaper", one has "longer zoom" and the other is "better for travel". This may help the user 1314 see alternatives that they may not have otherwise seen based on how they answered the questions 1320 leading up to the decision 1310.

[00109] In embodiments, users may be asked different types of questions, such as questions 1320 about the item being recommended (price, color, etc), questions 1320 about themselves, and the like. The system may differentiate users along dimensions, such as psychographic dimensions, demographic dimensions, and the like. Properties of users that may be predictive may include the user's age, sex, marital status, whether they live in rural/urban areas, frequencies of church attendance, political affiliation, aesthetic preferences, sense of irony/sense of humor, socio-economic background, taste, preference for neat or disorganized lifestyle, degree to which they plan ahead of time, and the like.

[00110] In embodiments, it may be difficult to directly ask questions 1320 and instead the system may try to measure things that are correlated instead. For example, instead of asking about income, the system might ask where the user 1314 prefers to shop (e.g., Wal-Mart, Target, Saks, etc). Aesthetics may be determined via showing pictures of art, living rooms, clothes, and the like, and asking which style the user 1314 prefers. In embodiments, pictures may take the place of the question (and the answers may be about how you react to the picture) or the picture can take the place of answers to questions 1320 such as "Which of the following best resembles the clothes you like to wear".

[00111] In embodiments, the system may group questions 1320 by whether they are about the item being recommended or about the user. The system may explain what type of questions 1320 it is asking in order to help the user 1314 understand the value of otherwise surprising and potentially offensive questions 1320 being asked. The system may also display other types of messages to the user 1314 while asking questions, such as telling the user 1314 how many questions 1320 remain, taunting the user 1314 by saying the system can already guess what decision 1310 to make, and the like.

[00112] In embodiments, instant messenger (IM) systems may provide a natural interface to the question 1320 and answer 1322 dialog of the system. For example, a user 1314 may invite our system to their "buddy list" and then initiate a dialog to get a decision 1310 over IM. The system may IM the first question 1320 to the user, the user 1314 may then IM their answer 1322 back, and the like, until eventually the system IM'ed the user 1314 a link to the decision, or directly IM'ed the name of the decision 1310 to the user. In embodiments, other forms of communications may also be used, such as cell phones, SMS, email, and the like.

[00113] In embodiments, the system, such as in the form of an application, may be embedded in third party web sites. For example, the system could be put on a website that sells cameras and offer to recommend relevant cameras to the user. Alternatively, after the user 1314 searched for cameras and had a list of potential cameras they were interested in, the system could ask questions 1320 to help the user 1314 decide amongst the list of cameras. For example, if all of the cameras that the user 1314 was considering were good for travel the system would not ask about how the user 1314 wanted to use the camera, but the system might realize that asking whether interchangeable lenses were desired could be used to recommend one camera over another.

[00114] In embodiments, the system may make decisions in a plurality of topic areas, such as: products (e.g., cameras, TVs, GPS/navigation, home audio, laptops, bath & beauty, baby, garden/outdoor, automobiles, jewelry, watches, apparel, shoes, and the like), travel (e.g., where to go, where to stay, what region to visit, what to do there, and the like), financial (e.g., which mortgage, whether to refinance, which credit card, whether something is deductible on taxes, what type of IRA to save in, asset allocation for investments, and the like), gifts for

various holidays and occasions, other date-based decisions (what to dress up for Halloween, and the like), personality (e.g., about a user's personality, about their relationships, their career, and the like), recommending the right pet, drinks and other aspects of night-life, books, movies, film, /music, concerts, TV shows, video games, where to eat, what to order, celebrity related such as which celebrity the user 1314 is most similar to, recommending a gift, what neighborhood to live in, what to watch on television, and the like.

[00115] In embodiments, the system may be used to diagnose problems, such as in the areas of technology / IT (e.g., computer, software, printers, home networking, wireless, business networks, performance issues, and the like), medical / health, automotive, relationship or interpersonal problems, home and building problems, and the like.

[00116] In embodiments, users of the system may be either anonymous or logged in users. A logged in user 1314 may be one that has created an account on the site. Logged in users may also have profile pages about them. Content on the profile page may include basic information about that user (nickname, picture, etc), decisions they have received and liked, decisions the system predicts the user 1314 will like even though the user 1314 has not answered questions 1320 in that topic area, lists of facts about the user 1314 that the user 1314 has given so that they do not need to be repeated each time the user 1314 uses the system for a decision (e.g., the user's age or their aesthetic preferences can be given once and remembered across different times the user uses the system), lists of tasks that the system thinks the user 1314 may be qualified and interested in doing via the wiki (such as reviewing new user 1314 submitted content, fixing spelling errors in user 1314 submitted content, reviewing new content found by the widgets, etc), other users with similar answers to questions, and the like.

[00117] In embodiments, users may also have various titles, ranks or levels which may affect what they can do on the system. For example, some users may be given the title of "moderator" in a particular topic which would allow those users to edit certain aspects of those topics. The ranks and titles may be assigned manually or by through automatic means including being based on how many decisions they have given, how many new questions 1320 or solutions they have contributed to the system, how many tasks they have

accomplished using the wiki, how well they answer 1322 certain questions 1320 in the various topics, and the like.

[00118] In embodiments, non logged in users may not have the benefit of using the system with a large selection of aesthetic or taste-based preferences already entered into their profiles. Based on learning or manual training from logged in users 1314, the system may select some aesthetic questions to ask in question dialogs when non-logged in users seek advice in particular topic areas. For example, based on logged in users answering taste questions about themselves and then giving feedback about which cars they like and don't like, the system may learn that a question 1320 about whether the user enjoys gourmet dining is useful to ask non-logged in users trying to decide between a Toyota and a Lexus. Using the attribute associations learned or manually specified by logged in users, the system may then adjust whether it recommends the Toyota or the Lexus to the non-logged in user.

[00119] In embodiments, the system may learn from users submitting feedback on decisions. Some users may either intentionally or unintentionally give incorrect feedback. For example, a vendor may try to game the system to make their product be highly recommended. Alternatively, a user 1314 who does not know much about video games may recommend a video game that in reality is not a good video game. The system may try to filter out feedback from these users by a variety of means. The system may throttle the number of feedbacks that a given user 1314 can submit (and have a higher throttle limit if the user 1314 is logged in or has a high rank/title). The system may also throttle or weight feedback based on how well the user 1314 answers certain 'test' questions 1320 during the question 1320 & answer 1322 phase in order to test the user's knowledge of the subject and weigh feedback from knowledgeable users more than unknowledgeable users. The system may also require the user 1314 to pass a 'captcha' (Completely Automated Public Turing test to tell Computers and Humans Apart) before their feedback is counted or they get a decision. The system may also look at the series of answers given by the user 1314 and weight the user's feedback based on that series of answers. For example, if the user 1314 either always clicked the first answer 1322 or the user 1314 clicked in a very improbable way, then the system may weight that user's feedback lower. Finally, the system may change the weight of the user's feedback or decide to not show a

decision 1310 based on the history of previous game plays. For example, the 10th time a user 1314 tries to get a camera decision 1310 the system may weight their feedback less than on the 9th time.

[00120] In embodiments, the system may include search engine optimization (SEO), the process of improving the system's website rankings within major search engines. This process may be broken down into several mostly-automated steps, such as discovering the keywords that users are searching for, understanding the competition in the search engines to have the site's page come up when users search for these words, understanding how search engines rank sites, understanding what changes to the system's website need to be made in order to increase the site's ranking for common searches, and the like.

[00121] In embodiments, discovering keywords that users may be searching for may be found through different means, such as using keyword suggestion tools such as what Google and Yahoo provide, using data about historical searches licensed from third party data providers and crawling other websites to see what words they use, and the like. Once these keywords are found, the system may use the data in many ways, such as bidding on those words via search engine marketing (SEM), developing content on the system's site about those keywords in hopes of getting search traffic in the future, looking at how our competitors are using those same keywords, and the like.

[00122] In embodiments, the system may understand what other sites are doing and how they rank in the search engines by running keywords through the search engines and looking at who is advertising on each keyword and what the top natural search results are for each keyword. The sites discovered through this process may be crawled to discover more potential keywords. The system may also decide to develop new content or avoid a market based on this competitive information. If there are few highly ranked sites in a content area, the system may develop content in that area.

[0100] In embodiments, the system may understand that paid advertisements that bring users 1314 to the site are relatively cheap in one topic area of advice on the site and expensive in another. The system may therefore try to advertise for the low-cost traffic, help those users 1314 with their decision, and then recommend that those users 1314 use the system in a topic area that is

expensive to advertise and buy traffic in. For example, the system may run ads for people who want to figure out what dog breed they should buy, help those users 1314 decide what dog breed is right for them, and then direct them to figure out where they should buy their pet medicines. The latter topic area being one that may be expensive for the system to source traffic in due to expensive ad rates, while the former topic area may be relatively cheap, as few existing businesses may be competing for customers who want advice on what type of dog to get.

[0101] In embodiments, the system may understand how search engines rank their natural (non-sponsored) search results by studying the relationship between sites that come up when a search is done and factors of those sites. Possible factors that may be correlated between sites that come up with high ranking may be factors such as the content of the site, number and quality of other sites linking to the site, the type of content on those other linking sites, and the like. From the prior step, the system may generate a list of site factors, ranked by their ability to increase a sites ranking in the search engines, and the like. The system may then use this ranked list to make changes to the site to increase the probability that the site as a whole, or certain pages on the site, will be highly ranked in the search engines.

[0102] Search engines may typically utilize a keyword index to find documents relevant to a user's query. In embodiments, the present invention may utilize a "decision index", which may also map user-input to relevant documents. The index may be built automatically, experts may hand build the index, the index may be learned through feedback from different types of users who implicitly or explicitly decide to train the system, and the like. The results of the search utilizing the decision index, may be displayed as a list of documents, a single document, and the like.

[0103] Referring to Fig. 1, an embodiment for a list of topics 102 in the system from which users may get decisions is presented, including cameras, cell phones, coffee and espresso, drinks, favorite celebrity, GPS devices, grills, Halloween, laptops, personality, toe rings, TVs, vacations, video games, watches, and the like. In addition, there may be an indicator as to the number of decisions learned 104 from user ratings, such as learned from 43,921 user ratings.

[0104] Referring to Fig. 2, an embodiment of an example question 1320 that the system may ask a user 1314 is provided. In this example the user 1314 is asking for a decision 1310 related to the purchase of a camera, and the question 1320 is "How much are you willing to spend?" The user 1314 may now choose from the selection 204, such as to select between less than \$200, up to \$300, up to \$500, more than \$500, I don't know, and the like. In addition, there may be an indication as to how many questions 1320 may be asked 202, such as in "In 10 questions or less, get camera decisions preferred by people like you." In embodiments, the user may also offer their own question, their own answer, their own decision, and the like, where the system may utilize this information in the current or future decision session. In embodiments, the user 1314 may choose to skip the question 208, where the user 1314 may now be provided an alternate decision based on a reduced amount of information available from the user, the system may ask the user alternate questions 1320 to make up for the skipped question 208, the question 1320 may have been a test question and will not affect the resulting decision 1310, and the like.

[0105] Referring to Fig. 3, an embodiment of an example picture question 1320 that the system may ask a user 1314. In this example, the system may be asking a question 1320 whose answer 1322 may better enable the system to determine a personal characteristic of the user 1314. For instance, the question 1320 as illustrated asks "Which of these causes you the most concern?", where the picture choices 304 are indicative of certain topics, such as pollution, finances, national defense, health, and the like. This question 1320 may be targeted to the current user or be inserted as an experimental question. In embodiments, the user 1314 may be informed that the question 1320 is an experimental question 302, such as shown in Fig. 3 with the header that reads, "Finally, please answer the experimental questions submitted by another user."

[0106] Referring to Fig. 4, an embodiment of an example of the type of information 402 the system may show the user 1314 when making a particular decision 1310 is presented. For example, the decision 1310 may be for a certain camera, where information is provided about the camera, such as a description, who uses it, the best cost for the camera, how it compares 404 to other cameras, and the like. In embodiments, other decisions 1310 may be provided, such as with a relative ranking 408, by a score, by a percentage matching, and the like.

The user 1314 may also be queried for feedback 1312, such as being asked if the decision 1310 is a good decision. In addition, the user 1314 may be provided with the opportunity to find out more about the decision 1310, such as more about the product 410, best price finder 412, websites to more advice, and the like.

[0107] Referring to Fig. 5 and Fig. 6, the user 1314 may be provided with various top lists 502 associated with a topic as described herein, such as presented in association with a decision, in association with a user's request to view top lists, and the like.

[0108] In embodiments, the present invention may provide users with a home page 700 including user 1314 identification 702, personal representation, past decisions made, future topics for consideration, decision 1310 to make today 714, and the like. Fig. 7 provides an example of a user home page 700, such as what the user 1314 sees when they are logged into the system account. Here, there may be a display of recent decisions the system recommended, lists of popular topics 708 to get decisions in, a search interface 710 to find topics, status updates about the user 1314 getting benefits for contributing to the system, recent activity 704, access to the user's profile 712, and the like.

[0109] Figs. 8 and 8A provide an example of a user's profile 712 page showing information about them and their account. The user 1314 may manage user information 802, such as a user's email address, password, and the like. They may also answer questions 1320 about themselves and have these answers remembered 810 and automatically used when they use decision making topics in the system. The user 1314 may also receive rewards 804, such as "badges", and see them displayed as received in response to helping other users, contributing to the system, and the like. Some of these rewards may be based on the quality of the user's contributions, on the quantity of contributions, and the like. In addition, users may be assigned a demographic group 808 of people who answered questions 1320 about themselves similarly.

[0110] In embodiments, users may be able to decide they want to contribute expertise 902 to the system, such as in a 'teach the system' mode. Fig. 9 shows an example of various links/pages that may allow a user 1314 to contribute, such as giving the system training about various decisions, rating the quality of pictures and user-contributed prose, finding duplicate items and

questions, contributing new decision making topics, contributing new questions 1320 to existing topics, and the like.

[0111] In embodiments, the user, after choosing a topic for the system to make decisions for, may be asked questions. Fig. 10 provides an example of how a question 1320 may be presented 1000 to the user. As shown, the presentation of the question 1320 to the user 1314 may provide different elements, such as a topic heading 1002, a picture or illustration associated with the topic 1004, a question, a set of answer choices, and the like.

[0112] After answering questions, the user 1314 may be provided an answer 1322 or decision 1310 associated with the user's original question. Fig. 11 and 11A show an example of how a decision 1310 may be presented 1100 to the user, and may include a primary decision, information summarizing the decision, alternate decisions, variations on the decision, and the like. In addition, the user 1314 may be provided with an opportunity to provide feedback 1312 to the system, such as whether the user 1314 agrees with the decision 1310 or not. The user 1314 may also be provided other suggested topics 1102, such as based on the current topic, answers provide, history of answers, a user's profile, a user's history of questions, topics that other users found helpful, and the like.

[0113] Fig. 12 shows an example list of decisions 1200 in a topic. For a product topic, such as shown, the "decisions" may be what product to buy. For other topics, the decision 1310 might be "yes, dump him" or "no, don't get a tattoo". The decisions may be ranked and ordered based on their relevancy to the user, based on how the user 1314 answered questions, based on how the user 1314 answered questions 1320 in the topic, and the like. Additionally, the items may be ranked by price, by name, and the like.

[0114] Fig. 16 shows an example of a contributor / expert interface home page 1600, showing recent contributions to the system 1602 and other users making contributions 1604. In upper right corner is a question for learning the user's taste preferences 1608.

[0115] Fig. 17 shows an example of a question in a dialog with the system asking an objective question 1700, and in this instance, to a user looking for help deciding what to name their new puppy.

[0116] Fig. 18 shows an example of a decision result showing the particular recommended decision 1800 (in this instance, name your dog Rusty),

reviews about this decision from other users 1802 (where it may be ranked, such as by their similarity to the user), yes/no buttons 1804 (such as for receiving feedback on this decision, showing other decision areas that the user might enjoy under, and the like), suggested Topics 1808. In this example, the system's second and third best recommended decisions are listed under the #2 tab 1810 and #3 tab 1812. The system may also be engaging in exploration by also recommending a "wild card" decision which may be a decision that was partly picked through randomness. The Suggested Topics 1808 may be selected based on how relevant the system thinks these topics may be for the user and/or how much profit the system thinks it may be able to generate from the user using these other decision areas.

[0117] Fig. 19 shows an example of an interface for users 1900 to set the associations between attributes and decision results. In this example, the decision result "Rusty" should be associated with the attribute "Is this name for a female, or male dog? Male".

[0118] Fig. 20 shows an example of how a user may edit content in the system 2000. In this example the user is able to edit a decision result: its name, description, URL for getting more information, etc.

[0119] Fig. 21 shows an example of how content that is editable by users may also have an interface for seeing prior revisions 2100 to the content and showing the changes between two prior revisions. Users may also revert the changes made by other users if those changes are deemed to be irrelevant or unhelpful. In this case the example shows the difference between two revisions to a decision result where the description of the result has been changed.

[0120] Fig. 22 shows an example showing a question being edited by a user 2200. New answers may be added, existing answers re-ordered, the question and answer text itself edited, and the like. Questions may be optionally "locked" to prevent other users from changing them, such as by indicated by the pad lock icon 2202.

[0121] Fig. 23 shows an example showing that edits to attributes may have revision histories like other editable content 2300. This example shows the difference between two revisions of the attribute associations between the decision result "Rusty" and the attributes "How many syllables do you want the name to have? No more than 2 or 3 or more is OK".

[0122] Fig. 24 shows an example showing a 'workshop' screen 2400 where newly added areas of advice may be first displayed. In embodiments, expert users may make additions here without regular users seeing the works-in-progress. Content that is deemed objectionable, irrelevant, or low quality may be voted on and removed from the system.

[0123] Fig. 25 shows an example showing the system asking the user taste / subjective questions 2500 in order to learn taste and subjective preferences from the user. After answering these questions the system may show statistics on how other users answered the same question.

[0124] Fig. 26 shows an example of an activity feed 2600 of recent activity by contributors across the site showing newly added content and expert training.

[0125] In embodiments, the present invention may provide a facility for providing an improved way to provide decisions to a user 1314 with a question 1320 across a broad category of topics, including products, personal, health, business, political, educational, entertainment, the environment, and the like. For example, the system may provide decisions on everything from whether a user 1314 should break up with their boyfriend, to whether you should get a tattoo or not, to whether you can deduct something on your taxes in addition to product decisions, and the like. In embodiments, the system may provide decisions on any interest a user 1314 may have.

[0126] In embodiments, the present invention may provide a decision system that is flexible and is capable of changing and growing. This may be partly enabled by the system's use of a dialog of questions 1320 and answers to make a decision, and then getting feedback from the user 1314 so the system can improve. In embodiments, this approach may be significantly more powerful since the system may ask any question 1320 and therefore get much better information from the user 1314 about their wants. In addition, users may be able to extend the system by entering their own questions 1320 and answers for the system to ask, entering in new decisions for the system to make, and the like. The system may then automatically try out newly entered information to see if it is useful or helpful and use this new information to determine if it is useful, and possibly stop asking/using the questions/decisions that may not be as helpful to

users. In embodiments, this approach may provide for building a wisdom-of-the-crowds based decision making expert system for potentially any topic.

[0127] In embodiments, the present invention may also provide improved decision facility to the user 1314 by providing decisions by ranking across non-traditional feature dimensions. For example, instead of just ranking cameras by price or size, the system may rank cameras based on how much they're liked by retired people, how sexy they look, and the like. The system may then help users navigate across these dimensions. For instance, instead of users just being able to say "I like this camera, but want a cheaper one" the system may let them do things like say "I like this camera but want one better for learning photography" or "I like this vacation, but want one with a more active social scene".

[0128] In embodiments, the present invention may lend itself to a variety of different user interfaces, such as a web interface, instant messaging, voice, cell phone, SMS/instant messaging, third party use (e.g. a widget on a third party, web service sold to a third party), and the like. For example, a voice interface may be well suited to the system since there may be a very limited vocabulary that the system must recognize, such as just the possible answers to each question. In this way, if the system can't understand a user response it may just move on to another question 1320 instead of annoying the user 1314 by asking them to repeat their answer over and over. In another example, the present invention may be integrated into a third party site, such as a search for a TV on an e-commerce website, where the present invention is a widget to help the user 1314 narrow down the results, or using the present invention as a widget in association with a real estate website to build an MLS query for the user 1314 to find a house that is a good match for them. In embodiments, the present invention may provide a user interface, both in regard to a physical interface and in the way questions, answers, and decisions are presented, that provides the user 1314 with a significantly improved way to obtain decisions on a great variety of topics.

[0129] In embodiments, the present invention may be integrated into third party products in such as way as to improve the third party's user interface and user satisfaction. For example some website services provide predictions through past purchase history. In this case, the present invention may be able to

explore a user's mood or intent, such as through asking explicit questions. In the case of search engines, the present invention may be able to detect when the user 1314 is trying to make a decision 1310 and then start to ask them follow on questions. In the case of forum sites, mailing lists, news groups, and the like, the present invention may provide improved access to decisions and decisions that were made by people similar to the user. For example, the present invention may be able to search through all the forum posts to find people who are in the same situation as the user, and providing what decision 1310 the forum community recommended to them.

[0130] In embodiments, the present invention may be able to extend e-commerce web application user interfaces. For instance, a user 1314 may start a product search with a keyword search and then ask questions 1320 to narrow down the results to the best decision 1310 for the user. The present invention may be able to provide a Q&A interface for picking a product once the user 1314 clicks into a category page. For example, after clicking cameras on the website, the user 1314 might see a first question. The present invention may be able to rank products along dimensions that are based on how users answer questions. For example, cameras might be ranked from best to worst 'travel camera' based on how people answer 1322 the question "What do you want a camera for?" Answer "Travel" and then whether they give positive or negative feedback to a particular camera. This may allow the e-commerce website to rank a list of camera keyword search results from best to worst travel cameras.

[0131] In embodiments, the present invention may be able to provide an improved search engine capability, such as detecting when a user 1314 is trying to make a decision 1310 and switching to a Q&A interface, based on the search results from a keyword search ask follow up questions 1320 to narrow down or re-rank the results, ask questions 1320 in order to build a keyword search query or to refine a search query, learn feedback based on which links a user 1314 clicks after being asked questions, and the like. In addition, the present invention may implicitly learn about the user 1314 and alter rankings based on these implicit facts, such as what time of day they're using the system, where they are in the world, what type of browser they're using, weather where they are, and the like.

[0132] In embodiments, the present invention may be able to provide a way for information to be gathered and utilized by users. For instance, Wikipedia is a way for users to contribute information such that the end user 1314 must, to some extent, self validate the accuracy of the information subsequently supplied to them. In a similar fashion, the present invention may be able to host a web application that utilizes user contributed content. For instance, instead of learning what the prices of cameras are, the web application could have users input the prices of cameras and then allow other users to self validate these claims. In this way, the scope of the contributed information may be allowed to grow organically as users interact with the system.

[0133] In embodiments, some e-commerce applications may provide for products and/or services that are associated with personal preference, and so may benefit from the present invention. For instance, there are currently several movie rental web services, where the user 1314 selects movies for delivery to their home through the mail. Decisions are also provided to the user 1314 based on what the user 1314 has selected in the past. However, choosing a movie may involve personal interests at the time of rental that cannot be determined by past selections, such as mood, intent, weather, are they going to be alone or with someone, their current personal relationships, and the like. These types of interests may be explored with the present invention through questioning, and as such, may provide a much more personalized match to the user's interests at the time of rental.

[0134] In embodiments, local search applications may be improved through the use of the present invention. For example, if a user 1314 wanted a decision 1310 on where to eat dinner, they might search for "dinner in new york" and find a website with suggestions targeted to the query. This interface however falls short when the user 1314 doesn't have a clear idea as to what keywords to include. For instance, the user 1314 might not know the key options for food and might not think to search for 'ethiopian food new york.' The present invention may have the advantage of being able to figure out what question 1320 it should ask in order to narrow down the possibilities. In embodiments, the present invention may be able to aid in the building of a search query.

[0135] In embodiments, the present invention may provide for an improved way to match up users and experts, users and other knowledge based

users, and the like. For instance, a service may be provided to collect users and experts on different topics. Users may then come to the web interface of the service and enter into a session of Q&A where the best match is determined. As a result of the questions, the system may provide a decision, where the profile of the expert or other user 1314 is provided, and where the user 1314 may be asked if they agree with the recommended individual. In embodiments, the user 1314 may be provided a home page where previous matches and communications may be kept, forwarded to friends, experts rated, and the like.

[0136] In embodiments, the present invention may provide a platform for a community based question 1320 and answer 1322 application. For instance, users may post questions 1320 to the system, and other users may be allowed to respond. In such a system, a user 1314 may receive answers from a single user, multiple users, an automated system, and the like, where the user 1314 may be able to choose which answer 1322 they feel is correct. This answer 1322 may be kept private, posted for others to view, posted as the correct answer, provided to the system, and the like. In embodiments, the system may use the questions 1320 and answers to further develop the system, provide more accurate answers to users, sort the answers provided to the user, filter the answers provided to the user, and the like. In addition, users of the system may provide feedback to answers provided by other users, contribute to filtering criteria for eliminating incorrect answers, and the like.

[0137] In embodiments, the present invention may be used as entertainment, through machine learning capabilities as described herein. For instance, a user 1314 may provide an input or think of an idea, such as a topic, a keyword, a category, a question, a feeling, and the like, and the system may make a guess as to what it is through a series of questions 1320 and answers. For example, the user 1314 may think of an object, such as baseball, and the system may utilize machine learning capabilities, such as geometric systems, to provide questions 1320 to the user. A typical question 1320 may relate to size, such as 'is it bigger than a toaster?' These questions 1320 may then be answered by the user, such as through multiple choice selection, fill in the blank, true/false, free response, and the like. The system may then continue the question 1320 and answer 1322 sequence until it has a guess, and provide this guess to the user. In embodiments, this process may continue for a fixed number of questions, a

random number of questions, a user 1314 specified number of questions, a system determined number of questions, a system specified number of questions, and the like. In embodiments, the system may provide the user 1314 with a user interface, such as through the Internet via a website, through a stand-alone computational device, through a mobile computational device, through a phone service, through a voice interface, in association with an instant messaging service, through text messaging, and the like. In embodiments, the system may be provided to a third party, such as a widget to another website, as an API to a third party application, and the like. In embodiments, the present invention may use non-neural networks for entertainment applications, such as playing games.

[0138] In embodiments, the present invention may provide a system to assist in the discovery of new drugs, where the system may provide an aid in the selection and combination of molecules in creating a new drug. For example, the system may ask the user 1314 about information associated with chemical parameters, such as solubility, reactivity, toxicity, and the like, and combine these with questions 1320 to probe the user's expertise in recognizing molecular structures. As the question 1320 and answer 1322 sequence progresses, the system may provide the user 1314 with insights as to which molecular structures may be stable and synthesizable. In embodiments, the process may continue until the user 1314 has an improved sense for what molecular combinations may make for a new drug, until the choice of new exploratory routes are available for presentation to the user, until a new potential drug is identified, and the like.

[0139] In embodiments, the present invention may provide for an image finder application, where the user 1314 may be assisted in identifying an image that fits some subjective criteria that is not necessarily explicitly known to the user. For example, a user 1314 may be involved in the development of a brochure for a company, where they have the text for the brochure, but need to select an image to support the ideas and emotions that the text is trying to convey. The user 1314 may in this instance have a subjective idea as to what type of photograph may be required, but not necessarily to the extent that they could specify a search with keywords. The user 1314 may instead first specify the source of the images, such as from a file, a database, a website service, from Google images, from an advertiser image bank, and the like. Then the user 1314 may be asked a series of questions, or be presented with a series of images to

choose from. The answers and/or selections that the user 1314 chooses may then be utilized in refining the choices that are next presented to the user, and from which further questions 1320 and/or image selections may be provided. In embodiments, this process may continue until the user 1314 finds an image to select as the final image. Additionally, the system may take the user's 'final selection' and select a group of other similar images for presentation to the user, at which time the user 1314 may choose to continue the process of selection refinement.

[0140] In embodiments, the present invention may be used in a baby naming application, where the user 1314 may have only a vague sense of what names they might prefer. The user 1314 may be initially asked different types of questions 1320 intended to provide the system with information to aid in the learning of the user's preferences, such as questions 1320 about family, friends, education, heritage, geographic location, place of birth, hobbies, books read, movies watched, and the like. The system may then continue to learn through the presentation of questions 1320 associated with name preference in a plurality of ways, such as rating name, choosing from a list of names, answering questions 1320 pertaining to name, and the like. In embodiments, this process may continue until the user 1314 finds a name to select as the final name. Additionally, the system may take the user's 'final selection' and select a group of other similar names for presentation to the user, at which time the user 1314 may choose to continue the process of selection refinement.

[0141] In embodiments, the present invention may provide decisions for a plurality of topics including, but not limited to, video games, laptops, vacations, cameras, general personality, drinks, cell phones, televisions, grills, watches, coffee machines, toe rings, Halloween, GPS devices, hottest celebrity, your personal hero, presidential election, baby toys, blogs, camcorders, cars, which star wars character are you, credit cards, hair care, skin care, sex and the city, should I get a tattoo, professions, how much allowance, city to live in, dog breeds, fragrance, New York, neighborhood chooser, software, desktop computers, DVD players and recorders, cigars, charities, Broadway shows, speakers, home theater systems, MP3 players, computer networking devices, headphones, memory cards, magazines, books, Oprah picks, books, The New York Times bestsellers, business casual clothing, franchises, cookware, toys,

toys - educational, athletic apparel, espresso machines, should I go Greek, should I come out to my parents, should I ask for a raise, do I have a drinking problem, should I medicate my add/ADHD child, vacuum cleaners, clothes washers and dryers, is working at a startup right for me, humidifiers, are you a good friend, risk of developing diabetes, which foreign language should I learn, microwaves, car audio, what kind of customer are you, wine, should I join the military, which military branch should I join, what kind of art will I enjoy, baby and toddler car seats, baby strollers, baby travel accessories, natural and organic beauty products, makeup, home audio receivers and amplifiers, copiers and fax machines, printers, breakup with my boyfriend/girlfriend, which Greek god are you, what game show would I enjoy, computer accessories, which superpower should you have, college, online degree programs, choose a major for college, identity theft prevention, should I hire a personal trainer, should I buy or lease a car, should I have laser eye surgery, what should I do about losing my hair, should I start my own business, should my child start kindergarten, how to entertain my family visiting NYC, OTC pain relievers, do I need a living will, miles or cash for my next flight, best way to whiten my teeth, should I let my daughter wear makeup, is hypnosis likely to cure my bad habit, ED options, sleep aids, OTC allergy pills, how much money to spend on a wedding gift, should I buy the extended warranty, is it better to take the SAT or ACT, personal audio accessories, coffee/espresso drink would I enjoy, video game consoles, jeans, downloadable PC games, snacks, vitamins and supplements, which superhero am I, sunglasses, kitchen gadgets, pillows, beauty accessories, beauty bags and cases, sporting goods, which musical instrument is right for me, should I hire a decorator, electronic readers, where do you belong in a shopping mall, power washers, small business, phone system, how much to tip, should I try Botox, should I get liposuction, risk of skin cancer, should I refinance my home, car services (NYC), microbrewery beer, gourmet chocolates, am I saving enough for retirement, entertainment centers/TV stands, cookbooks, electric shavers, keep sending nieces/nephews bday gifts, luggage, computer projectors, energy/workout bars, razors, gourmet ice creams, online dating, newscasts, makeup, tools and brushes, beauty mirrors and compacts, business books, how soon to call after a first date, places to retire, external hard drives, universal remote controls, walking shoes, should I sell my life insurance policy, how green

are you, do I have an eating disorder, baby cribs, diets and diet books, cell phone plans, wedding and engagement rings, am I assertive enough, does my child play video games too much, tax preparation (personal return), should I get a reverse mortgage, cancel plans with friends for a date, children's TV shows, kitchen countertops, bathing supplies, insect repellents, cancer specialist, hospitals, national chain restaurants, cereal, should I have kids now, should I hire a nanny, movies, beef cuts, target calories per day, do I have OCD, home air purifiers, auto air fresheners and purifiers, i-phone applications, gay/lesbian vacations, is it ok to ask my co-worker on a date, is my pre-teen ready to babysit, sports/energy drinks, TV shows, office furniture, motorcycles, reward a child for a good report card, lawn trimmers and edgers, am I too stressed, religion, do you make a good first impression, do you spend too much time online, should I get a new hairstyle, should I home-school my child, diaper bags, should I use cloth or disposable diapers, dog toys, is my partner cheating on me, classic books should my elderly parent stop driving, am I over my ex, is it lust or love, pedometers and heart rate monitors, chewing gum, weather devices, will gas additives help my car, Orlando theme parks, how big of a turkey should I buy, popular music - new releases, self tanner, tax and money management, software, baby bottles and Sippy cups, baby high chairs and booster seats, baby tethers, toasters and toaster ovens, comforters sheets and bed linens, flatware sets, pet carriers and kennels, cheese, kitchen faucets, casual shoes, dress shoes, beauty electronics, am I saving enough for retirement, mutual fund chooser, steak cuts, what is my D&D alignment, acne and pimple medication, bathroom faucets, home exterior lighting, landscape lighting, lawn mowers, aperitif, cognac, gin, rum, scotch, tequila vodkas, whiskeys, Las Vegas shows, sunscreen, running shoes, US MBA programs, patio and outdoor furniture, kitchen knives, are you a true fan, auto insurance, personal legal services, should I hire a financial advisor, indoor plant selector, delivery services, can I deduct it, pool heaters, sofas, house numbers, contact lenses, birthday gifts, has my career peaked, electronic books, doorknobs & lock sets, snow removal equipment, green home improvement, kids clothing & swimwear, motorcycle helmets, bicycle helmets, juicers, golf clubs, refrigerators, wine coolers, ranges and ovens, air conditioners, Christmas gifts, breakup phrases, cold sore medication, diabetes monitoring devices, smoking cessation, what do I do about the hair on my back, hormones to counteract

menopause, hiking backpacks, school backpacks, get a website/domain, e-mail services, web hosting, carpets, power tools, tile, water heaters, outdoor paint, window treatments, fireplace screens, indoor lamps, small business legal services, brunch recipes, ceiling fans, mattresses, Las Vegas hotels and casinos, salsas, love quiz for valentines, how much to spend on client gifts, anniversary gifts, outdoors outerwear, casual outerwear, camping tents, sleeping bags, tires, adventure vacations, music downloads, video downloads, wedding dresses, wedding themes, Manhattan gyms, budget hotel chains, golf courses, ski vacations, US spas, ETF funds, designer handbags, should I declare bankruptcy, 401k as down payment on home, should I see a psychiatrist, self defense, dishware, dishwashers, political parties, new year's resolutions, cruise lines, family vacations, baby food, baby health care products, should I shave my head, t-shirts, online photo services, buy a class graduation ring, summer job/internship, where to volunteer, home alarm systems, diagnose your relationship issues, is she/he hot for me, should I adopt, should my aging parents be driving, online bank accounts, BBQ sauces, frozen pizza, recipe finder, should I re-gift it, bodybuilding supplements, home workout equipment, how many hours of sleep do I need, should I consider plastic surgery, risk of arthritis, risk of heart disease, risk of osteoporosis, do I have a gambling problem, best dance to learn, bicycles, cat food, dog food, hobby recommender, martial arts, poster art, outdoor flower selector, which Muppet are you, activities for kids, how ethical are you, should I baptize my child, Miami hotels, US national parks, motor oils, automotive video, blouses, coats, dresses, glasses frames, hosiery, interview clothes, jackets, negligee, pants, shirts, skirts, hats, phones-land lines, steakhouses, which birth method is right for you, summer camp recommender, march madness bracket chooser, baby formula, New York bakeries, fractional jet ownership, how self confident am I, digital photo frames, do I need an accountant, does my child have ADD/ADHD, document shredders, baby monitors, green home improvement, conference phones, and the like.

[0142] In embodiments, and as depicted in Fig. 13, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1302. The process may begin with an initial question 1320 being received 1304 by the machine learning facility 1318 from the user 1314. The user 1314 may then be provided with a dialog 1308 consisting of questions

1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide a decision 1310 to the user based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, future questions 1320 and decisions 1310 provided by the machine learning facility 1318 may be improved through feedback 1312 provided by the user 1314.

[0143] In embodiments, the initial question 1304 posed by the user 1314 may be an objective question, a subjective question, and the like. A question 1320 may be provided from amongst a broad category of topics, such as topics pertaining to a product, personal information, personal health, economic health, business, politics, education, entertainment, the environment, and the like. The questions 1320 may be in the form of a multiple choice question, a yes-no question, a rating, a choice of images, a personal question, and the like. The questions 1320 may be about the user 1314, provided by another user, provided by an expert, and the like. The questions 1320 may be based on a previous answer, such as from the current dialog 1308 with the user 1314, from a stored previous dialog 1308 with the user 1314, from a stored previous dialog 1308 with another user. The question 1320 may be a pseudo random question, such as a test question, an exploration question 1320 that helps select a pseudo random decision 1310 on the chance that the pseudo random decision 1310 turns out to be useful, and the like. The questions 1320 may include at least one image as part of the question. The questions 1320 may be along psychographic dimensions. In embodiments, the questions 1320 may not be asked directly to the user 1314, but rather determined from contextual information, such as through an IP address, the location of the user, the weather at the user's location, a domain name, related to path information, related to a recent download, related to a recent network access, related to a recent file access, and the like.

[0144] In embodiments, the dialog 1308 may continue until the machine learning facility 1318 develops a high confidence in a reduced set of decisions, such as a reduced set of decisions presented to the user, a single decision 1310 presented to the user. The decision 1310 provided by the machine learning facility 1318 may be independent of the order of questions in the dialog 1308. The decision 1310 may provide an alternate decision 1310 when at least one

question 1320 in the dialog is omitted, where the alternate decision 1310 may be different based on the machine learning facility 1318 having less information from the user 1314. The decision 1310 may display a ranking of decision choices, such as ranking decisions across non-traditional feature dimensions. The decision 1310 may display at least one image related to the decision 1310. The decision 1310 may be a pseudo random decision on the chance that the pseudo random decision 1310 turns out to be useful, such as the pseudo random decision being part of a system of exploration, where the system of exploration may improve the effectiveness of the system, the machine learning facility 1318 may learn from exploration, and the like.

[0145] In embodiments, the feedback 1312 provided may be related to, or derived from, how the user 1314 answers questions 1320 in the dialog 1308, how the user 1314 responds to the decision 1310 provided by the machine learning facility 1318, and the like. In embodiments, the feedback 1312 may be solicited from the user 1314.

[0146] In embodiments, users 1314 may extend the learning of the machine learning facility 1318 by entering new information, where the new information may be their own topic, question, answer, decision, and the like. The machine learning facility 1318 may use the new information to determine whether the new information is helpful to users.

[0147] In embodiments, a user interface may be provided for user interaction with the machine learning facility 1318, such as associated with a web interface, instant messaging, a voice interface, a cell phone, with SMS, and the like.

[0148] In embodiments, the present invention may help a user make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314, where the initial question 1304 may be associated with one of a broad category of topics, such as product, personal, health, business, political, educational, entertainment, environment, and the like. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question

1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, future questions 1320 and decisions 1310 provided by the machine learning facility 1318 may be improved through feedback 1312 provided by the user 1314.

[0149] In embodiments, and as depicted in Fig. 14, the present invention may help a user make a decision 1310 through the use of a computing facility 1402. The process may begin with an initial question 1304 being received by the computing facility 1418 from the user 1314. The user 1314 may then be provided with a dialog 1408 consisting of questions 1320 from the computing facility 1418 and answers 1322 provided by the user 1314. The computing facility 1418 may then provide a decision 1310 to the user 1314 based on an aggregated feedback 1428 from the feedback from a plurality of users 1412. In embodiments, the computer facility 1418 may improve future questions 1320 and decisions 1310 provided by the computing facility 1418 based on receiving feedback 1412 from the user.

[0150] In embodiments, the present invention may help a user make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314, where the number of questions 1320 and answers 1322 provided through the dialog 1308 may determine the quality of the decision 1310. The machine learning facility 1318 may then provide a decision 1310 to the user based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, future questions 1320 and decisions 1310 provided by the machine learning facility 1318 may be improved through feedback 1312 provided by the user. In embodiments, the quality may be high when the number of questions 1320 and answers 1322 large, such as greater than 10 questions, greater than 15 questions, greater than 10 questions, and the like. In embodiments, the quality may be good quality when the number of questions 1320 and answers 1322 is small, such as less than 10 questions, less than 5 questions, less than 3 questions, one question, and the like.

[0151] In embodiments, and as depicted in Fig. 15, the present invention may make a decision 1310 through the use of a machine learning facility 1318. The system may include a machine learning facility 1318 that may receive an initial question 1304 from the user 1314, a dialog facility 1502 within the machine learning facility 1318 providing the user 1314 with questions 1320 and accepting answers 1322 from the user, the machine learning facility 1318 providing a decision 1310 from a decision facility 1504 to the user 1314, and the like. In embodiments, the decision 1310 provided to the user 1314 may be based on the exchange of dialog 1308 between the user 1314 and the machine learning facility 1318, and pertain to the initial question 1304. Further, the machine learning facility 1318 may receive feedback 1312 through a feedback facility 1508 from the user 1314 to improve future questions 1320 and decisions 1310 provided by the machine learning facility 1318.

[0152] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314 through a third party, such as a search application, a social network application, a service provider, a comparison shopping engine, a media company's web environment, and the like. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, future questions 1320 and decisions 1310 provided by the machine learning facility 1318 may be improved through feedback 1312 provided by the user 1314.

[0153] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314 through a third party search application, where the user 1314 begins with a keyword search on the third party search application and then is provided a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the

user 1314. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, where the decision 1310 may be provided back to the third party search application, such as in the form of a sorted list.

[0154] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314, where the machine learning facility 1318 may utilize third party information, functions, utilities, and the like. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, third party information, functions, utilities, and the like, may include an application programming interface (API) enabling the collection of cost information, product information, personal information, topical information, and the like.

[0155] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314 through a third party search application, where the user 1314 begins with a keyword search on the third party search application and then is provided a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, the decision 1310 may be provided back to the third party search application based at least in part on collaborative filtering.

[0156] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314. The user 1314 may then be provided

with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide at least one image with the decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, the image may be a photograph, a drawing, a video image, an advertisement, and the like.

[0157] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314 where the questions 1320 may be determined at least in part from learning from other users of the machine learning facility 1318. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, the decision 1310 may be based at least in part on learning from decisions 1310 provided by other users of the machine learning facility 1318.

[0158] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, the decision 1310 may be based at least in part on collaborative filtering.

[0159] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine

learning facility 1318 from the user 1314. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, the decision 1310 may be based at least in part on collaborative filtering whose context is provided through the dialog 1308, such as at least one question providing the context for the collaborative filtering.

[0160] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, the decision 1310 may be based only on information gathered through a plurality of user 1314s of the machine learning facility 1318 and pertaining to the initial question 1304, where at least one of the plurality of user 1314s of the machine learning facility 1318 may be the user 1314 associated with the dialog 1308.

[0161] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, and with limited initial machine learning facility 1318 knowledge on the subject matter of the initial question 1304. In embodiments, the limited initial machine learning facility 1318 knowledge may be seed knowledge, may be limited to basic knowledge

associated with the subject matter of the initial question 1304, may be limited to basic knowledge associated with the subject matter of the initial question 1304 where the basic knowledge may be expert knowledge.

[0162] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like, where the decision 1310 may be based on learning from a combination of expert and user inputs.

[0163] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide a category-based decision 1310 to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like.

[0164] In embodiments, the present invention may help a user 1314 make a decision 1310 through the use of a machine learning facility 1318. The process may begin with an initial question 1304 being received by the machine learning facility 1318 from the user 1314. The user 1314 may then be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314. The machine learning facility 1318 may then provide a decision 1310 to the user 1314 where the machine learning facility 1318 may utilize responses from a plurality of user 1314s of the machine learning facility 1318 to categorize and provide decisions 1310 along at least one of psychographic and demographic dimensions.

[0165] In embodiments, the present invention may provide a user 1314 with a response through the use of a machine learning facility 1318. The user 1314 may be provided with a dialog 1308 consisting of questions 1320 from the machine learning facility 1318 and answers 1322 provided by the user 1314, where the questions 1320 from the machine learning facility 1318 may be related to an application, such as an entertainment application, a drug discovery application, a baby name application, and the like. The machine learning facility 1318 may then provide the response to the user 1314 based on the dialog 1308 and pertaining to the initial question 1304, such as a recommendation, a diagnosis, a conclusion, advice, and the like. In embodiments, future questions 1320 and decisions 1310 provided by the machine learning facility 1318 may be improved through feedback 1312 provided by the user 1314.

[0166] In embodiments, the present invention may provide results based on multiple dimensions, such as a result based on a textual match from a user input, based on the user's taste profile, and the like. Fig. 27 shows an example search interface showing search results for the user query "fios". In this example, the rankings are based on first finding decisions and decision results (recommendations) that are a good textual match for the user query and then second ranking the decision results by the knowledge the system has about the user's taste profile. In this example, "fios" is a good textual match for the decision results "Verizon FIOS" as the recommendation to "Which ISP should I use" 2702 and "What US satellite/cable service provider should I get" 2704 and both are this users #1 ranked result based on the user's taste preferences.

[0167] In embodiments, the present invention may provide multiple question and answer 'results' in response to a user's unstructured input. For example, Fig. 28 shows a case where the user's query is 'suv'. The first result is for the topic 'What new car should I buy?' 2802 but with the question 'what type of car do you want?' already answered with 'suv'. This may provide a bridge between unstructured search and the structured Q&A data that the system stores. Further, as shown in the example, the user's top 3 results are displayed personalized based on their taste profiles. Effectively, the user has done a keyword search and gotten results without explicitly answering any questions via the traditional Q&A interface, such as for the other questions provided as

examples shown, 'Which BMW should I buy?' 2804, 'What late model used car should I buy?' 2808, and 'What tires should I use on my car or truck?' 2810.

[0168] In embodiments, the present invention may ask the user to express their preferences when they have given answers to questions that either contradict each other, are mutually exclusive, or which each individually have a dramatic affect on the rankings of results for the user. For example, the user may start the "What new car should I buy" topic and answered that they want a SUV that is under \$18,000 and is more practical than extravagant. The system may want to get the user's preference as to what is more important - that the vehicle be an SUV or under \$18,000. Fig. 29 shows an example question to the user, asking for their preference.

[0169] In embodiments, the system may learn a set of question importances that are per-user, per-decision-result, per-question's-answer-per-decision-result, and the like. For example, the system may learn that user A cares more about weight than price when it comes to buying small cameras, but cares more about price than anything else when it comes to buying cars.

[0170] In embodiments, users may extend the system by adding new decision results. Those new decision results may optionally include links to web pages to read more about the decision result. The system may automatically convert those links to affiliate links such that the system receives commissions from the site that the link points to. Further, based on the link submitted by the user, the system may recognize what kind of link it is and understand how to parse out information such as price for products, such as from Amazon.com, or parse out product codes so that vendor-specific API calls can be made to look up product information based on the product codes.

[0171] In embodiments, users may be clustered into groups using dimension reduction techniques such as singular value decomposition (SVD), eigenvector, and other like based approaches. The system may display information about why a group of users were clustered together. One way to do this is to find the top X dimensions in the low-dimension space that the cluster as a whole differs most on from the population average. The divergence of the cluster's distribution of answers from the general population's distribution in each dimension in the sub-space may be used to rank dimensions in terms of how well they explain what is unique about each cluster.

[0172] In embodiments, the dimensions in the sub-space may not be easily described or interpreted due to their being composed of many different features such as how people answered questions or which decision results they like. One way to explain what each dimension means may be to find questions and answers that are most correlated with the different extremes of the dimension and label the dimension with these questions/answers.

[0173] In embodiments, one way to cluster users may be to pick an initial random grouping of users and iteratively move users between clusters to minimize how much users differ from each other in their own cluster. After some number of iterations the process may be stopped or the process may continue until a threshold amount of error has been reached.

[0174] In embodiments, the present invention may facilitate matching by user similarity. Given a username, email address, numeric user id, and the like, provide a list of other users who are similar or dissimilar in some way. For example, given a Facebook username provide a ranked list of other Facebook users who have similar tastes over all or in some specific way, such as in electronics. In addition, this list may be optionally restricted, such as to other users being one degree away in the social graph from the first user (e.g. rank my friends on Facebook by their similarity to me so that I can ask the one most like me a question). In embodiments, user similarity may be computed via asking a user questions about themselves, looking at their social graph, using context like their location, IP address, time, and the like. The social graph may be used by mapping users in the social graph to known users in other data sets based on heuristics on username, email address, first name, last name, birthday, address, sex, and other like information. Adjacent people may be searched in the social graph to bring more people into consideration even if they are more removed from the person you are trying to 'triangulate' in on. For example, combining a social graph from Facebook with users who have written reviews on Amazon to find users on Facebook who are most like me and then look at what laptops they tend to like on Amazon in order to give me a laptop recommendation. In another example, reviews may be filtered on a site such as Yelp, Tripadvisor, Amazon, and the like, based on people similar to you. This information may then be used to help the user, such as recommending users to "friend" on Facebook, to 'follow'

on Twitter, and the like. Figs. 30 and 31 provide examples of similarity profiles as may be provided by the present invention.

[0175] In embodiments, the present invention may provide a level of indirection between recommendations by instead recommending people who in turn like/dislike things. The present invention may recommend things to buy based on what similar users bought on Amazon, places to eat based on similar users on Yelp, Zagat, Foursquare, and the like, things to click on in Google search results based on what similar people clicked on, and the like. For example, consider the problem of recommending which users a new user to Twitter should follow. The present invention may look at all the users on Twitter and who follows them, and match some of those followers to the data set of the present invention knowing things about each user based on the questions they have answered about themselves. This could also match the followers to other users, such as Amazon users, Yelp users, and the like, to learn other things about them. Based on this, the present invention may make inferences about the followers of a particular user, such as a Twitter user. Now, a new user may be asked about themselves and find which existing user's followers this new user is most like. A recommendation may then be generated for the new Twitter user to follow the existing Twitter user whose followers are most like the new user.

[0176] In embodiments, the present invention may facilitate real-time personalization, such as making recommendations that immediately reflect new information from a user, immediately use their social graph, new facts about them, their context such as changing location etc. to re-rank recommendations or otherwise improve results.

[0177] In embodiments, the present invention may match based on context, such as location, time, weather, social graph, and the like, such as implicitly using location to show nearby places a user might like to eat, drink, sites they might like to see, things to do, etc. For example, this process may then be used in a mobile application that has access to location data via GPS. Optionally, the recommendation may be informed by other parts of the user's context, such as the current weather (e.g. don't recommend a place that people only like because of the patio if it's raining), current time (e.g. don't recommend nightclubs at 10am), calendar (e.g. know when user is busy and where they have

to be in the future), social graph (e.g. recommend places that the user's friends are at all else being equal), and the like.

[0178] In embodiments, the present invention may provide for a natural language question and answer interface, such as to allow freeform or structured input from a user about a decision or recommendation they want help with. The input may be categorized either by asking the user, having other users review the question, using automated techniques like natural language processing (i.e., "is this question about electronics, travel, cars or some other topic?"), and the like.

[0179] In embodiments, the present invention may find similar users to the requesting user who has expertise in the category of the question. For example, the user may want advice about hotels in L.A. from people who have similar taste in hotels AND know something about L.A. hotels (such as either self-described knowledge or demonstrated knowledge based on their actions). For example, the present invention may then alert those similar users to the new question and requesting their help in solving it. Consideration may be given to how many prior questions/alerts they have been sent, how many they have already responded to, how helpful their responses have been, and the like. Similar users may be allowed to engage in a dialog with the requesting user to help inform the requesting user's decision or recommendation problem. The resulting dialog may be stored for others to use and encourage the similar users to index the dialog into a structured form to aid later recall by other requesting users.

[0180] The present invention may provide third-party sites with the ability to learn about their users, finding similar users, and making recommendations, such as independently of direct user interaction with the system. In embodiments, the present invention may provide a taste and preference API that third parties, such as third parties hosting their own web sites, may use to learn about the tastes, preferences, likes, dislikes, and other attributes of a user, where the user is not engaged in a dialog or interacting directly with the computing facility of the present invention. For example, a user may go to a website such as Amazon.com and make an inquiry about a product. In this instance, Amazon may have an API of the present invention that enables the creation or enhancement of a taste and preference profile of a user so as to better determine the user's tastes, preferences, and the like, and so enable the

third-party to better target meaningful responses back to the user with regard to the user inquiry. In embodiments, the third-party may use the API to determine the tastes, preferences, and the like of the user without user interaction, such as by determining their tastes and preferences through previous interactions with the user, where these previous interactions may be from previous interactions while on the third-party site, such as Amazon in this example, or from previous interactions with other websites hosting a similar API or through direct interaction with a facility of the present invention. In embodiments, the API may be used by a plurality of users such that interactions with the users may be used to identify other similar users, and so use the choices, decisions, selection, recommendations, and the like, of these other similar users to aid in the selection of recommendations to the present user. In embodiments, these other similar users may be associated with the third-party hosting the API, or from another host API site, or from a facility of the present invention. In embodiments, the use of an API of the present invention hosted by a third-party site may provide a significant advantage to the third-party site with regard to recommendations that are relevant to the tastes, preferences, likes, dislikes, attributes, and the like, of the user.

[0181] In embodiments, tastes and preferences of a user may be determined or augmented through other users, such as other similar users, other users that are connected to the user in a social network, other users that are associated through a personal or professional activity, other users that are friends or family, and the like. In embodiments, this may be done without the need to ask questions of the user. For instance, a user may have an existing taste and preference profile as established through the present invention, and that profile may be improved by collecting or inferring information about other users in their social network, family, place of business, and the like. In embodiments, the user's profile may be improved through inferring additional tastes and preferences from other similar users, or users shown to have some connection to the user, such as through a social network. In embodiments, the user's profile may be improved through choices made by other similar users, such as in a product selection, recommendation, and the like. In embodiments, the system may learn about a known user's taste profile through their rating things they like

and don't like, or through the use of natural language processing, such as inferring a taste profile by analyzing how the user tags their user profile.

[0182] Referring to Fig. 32, in embodiments the present invention may provide for a computer program product embodied in a computer readable medium that, when executing on one or more computers, helps a third-party website to learn about a user through the use of a computer facility 3202 by performing the steps of: (1) providing a user preference learning API to the third-party website 3220 to determine preferences of the user 3218 as applied to a market of the third-party, wherein the preference learning API is executing as an extension of the computer facility 3204; (2) receiving third-party information related to the market of the third-party 3208; (3) collecting the preferences of the user 3218 and storing them as a user preference profile 3210; (4) receiving a query from the user at the third-party website 3220 associated with the market of the third-party 3212; and (5) supplying a recommendation to the third-party based on the user preference profile and the third-party information to aid the third-party in the answering of the received query 3214. In embodiments, the determining preferences may be through the use of natural language processing. The computing facility may be a machine learning facility. The third-party information may consist of at least one of product information from product manufacturers, product information from web merchants, pricing information from other websites, availability information from other websites, pricing information from merchants, availability information from merchants, a review, comments, and ratings. The preference learning API may enable the collection of at least one of cost information, product information, personal information, and topical information. Further, user profile preferences may be additionally based on information inferred from a user's social network, where the user may not receive additional dialog between the user and the computer facility.

[0183] Referring to Fig. 33, the present invention may provide for the use of a taste and preference API to target responses returned to a user, such as for targeting advertising, show reviews from similar users, recommend products or services, show similar people on social networks, to rank search results based on which results similar users clicked on most, and the like. In embodiments, the present invention may provide for a computer program product embodied in a computer readable medium that, when executing on one or more computers,

helps target responses returned to a user through the use of a computer facility 3302 by performing the steps of: (1) providing a user preference learning API to a third-party website 3320 to determine preferences of a user as related to a market of a third-party, wherein the preference learning API is executing as an extension of the computer facility 3304; (2) receiving third-party information related to the market of the third-party 3308; (3) collecting the preferences of the user 3318 and storing them as a user preference profile 3310; (4) receiving a query from the user 3318 at the third-party website 3312; and (5) using in the user preference learning API the third-party information related to the market of the third-party, and the preferences of the user 3318 as stored in the user preference profile to provide a response back to the user that relates to the query from the user 3314. In embodiments, the computing facility may be a machine learning facility.

[0184] The response may be providing an advertisement to the user, where the advertisement may be based on the preferences of the user as stored in the user preference profile. The advertisement is provided by the computer facility, provided through the third-party and enabled through preferences provided to the third-party from the computer facility, delivered to other users in a user's social network, and the like. The response may provide a recommendation of a product, service, and the like related to the market of the third-party. Collecting the preferences of at least a second user may form a user preference profile for the second user, determining the second user to be similar to the user based on a comparison of preference profiles. The response may be providing a recommendation made by the second user. The collecting of preferences for the second user may be taken from an internet based social construct, and the response provides information to the user that shows the second user as a similar person on the construct, where the internet based social construct may be a social network. The collecting of preferences for the at least second user may include search result selections, the query may be a search request, and the response may be a search result ranked according to the search result selections of the at least second user. The collecting may be from recommendations, purchases, and search result choices made by the user. The collecting may be from sources that reveal location behaviors of the user. The source may be a user location information, such as from the web services

foursquare, yelp, Google, Gowalla, Facebook, and the like. The source may be user location information from a service provider. The third-party information may consist of at least one of product information from product manufacturers, product information from web merchants, pricing information from other websites, availability information from other websites, pricing information from merchants, availability information from merchants, a review, comments, and ratings. The determining preferences may be through the use of natural language processing. The API may enable the collection of cost information, product information, personal information, topical information, and the like. The collecting may be from a user's interactions as represented in an Internet based social interactive construct, where the internet based social interactive construct may be a social network. The response may be a list of reviews sorted by a review author's similarity to the user reading the reviews. Collecting the preferences of the user may be through third party websites. The collecting may be through crawling third party websites.

[0185] In embodiments, the present invention may utilize a taste and preference API that third parties may use to target advertising to a user based on the user's preferences, where the user is not engaged in a dialog or interacting directly with the computing facility of the present invention. For instance, a user may have previously undergone a dialog of questions and answers with the system, and through that dialog the system may have developed a taste and preference profile for the user. In embodiments, the dialog may have been provided directly with a facility of the present invention, or through a third-party API provided by the present invention. Alternately, the user may have never interacted with a facility of the present invention, where the user's taste and preference profile may be created and updated through the user's interactions, responses, recommendations, reviews, and the like. In embodiments, the system may learn about a known user's taste profile through their rating things they like and don't like, or through the use of natural language processing, such as inferring a taste profile by analyzing how the user tags their user profile. The taste and preference profile for the user may then be used to target advertising to the user, such as advertising that is matched to the user's tastes and preferences. For example, a third-party taste and preference API may be associated with an outdoor store website, such as L.L. Bean, REI, EMS, and the like, where the

outdoor store is trying to improve their targeting of advertising to their customers. A customer may then visit the outdoor store website and make a query for a product, such as for hiking boots. The taste and preference API may then enable a look up of the user's tastes and preference in order to establish a match for an advertisement placement to the user's browser. In this example, the user's taste and preference profile may indicate that the user enjoys traveling to New England, likes to camp, has a family with children, and the like. As a result, the advertisement facility associated with the website may select an advertisement that utilizes the information in the user query, in this instance hiking boots, and information from their taste and preference profile. An advertisement in this case may be for lodging in the White Mountains, which combines the user's hiking boot query with their preference for travel in New England. Further, the lodging may be a family lodging because of the user likes to travel as a family, and with attributes that mirror the user's preferences. In embodiments, the third-party taste and preference API may enable a third-party to improve their targeting of advertisements to users such that they are able to increase revenue made on a given advertisement placement from an advertisement sponsor. In embodiments, the user taste and preference profile may be developed in real-time as focused on their immediate query. Tastes and preferences may be gleaned so as to better target advertisements to the user, such as during subsequent refinement of product search, at the point of purchase, and the like.

[0186] In embodiments, advertising may be targeted to a user or a group of individuals associated with the user based on taste and preferences inferred through the user's social network. For instance, a taste and preference API used by a third-party may be used to establish taste and preferences for a group, node cluster, and the like, such as in a user's social network. In embodiments, the tastes and preferences inferred from the social network may utilize taste and preference profiles previously formed, such as through third party sites or through a facility directly associated with the present invention. These taste and preferences may then be used to better target advertisements to the user or to members of the user's social network. In an example, a third-party may want to target advertisement to a user, where the user has an established taste and preference profile stored in a facility of the present invention. The third-party may then use information in the user's profile to target advertisements. Alternately, the third-party may additionally use information inferred from a social network that the user is part of, such as information pertaining to the topics of the social network, common interests of users associated with the user in the social network, and the like. For instance, the user may have a taste and preference profile that indicates they are middle aged, politically conservative, rural, and the like, and is associated with users in a social network that have hunting as a dominate interest. In this instance, the third-party may target advertisements for hunting equipment, hunting trips, and the like, where the advertisement has been selected based on the user's existing taste and preference profile plus inferences from the user's social network. In embodiments, this may be done without the need to engage the user in a dialog as described herein, but rather indirectly through the user's interactions, such as on the third-party website, through third-party supplied information, on other websites where the present invention has a third-party API, and the like. In embodiments, the third-party may also utilize the taste and preferences from the user to target advertisements to other members of the user's social network.

[0187] In embodiments, advertising to a user may be targeted based on the product choices, recommendations, and the like, of users with similar tastes and preferences. For example, a first user may have similar tastes and preferences to a second user, where the first user has an existing taste and preference profile and has made certain product choices, recommendations, and

the like. Advertisements may then be targeted to the second user based on the decisions of the first user. For example, a first user may have a profile that indicates they are older, retired, lives in California, enjoys travel, and the like, where they have previously made a product choice for luggage. A second user may then be provided a recommendation for similar luggage based on the similarity of the tastes and preferences of two users. In embodiments, this may be executed without the need to engage either of the users in a dialog.

[0188] In embodiments, the present invention may provide for a computer program product embodied in a computer readable medium that, when executing on one or more computers, helps target advertising to a user through the use of a computer facility by performing the steps of: (1) providing a user preference learning API to a third-party website to determine preferences of the user as applied to a market of a third-party, wherein the preference learning API is executing as an extension of the computer facility; (2) receiving third-party information related to the market of the third-party; (3) collecting the preferences of the user and store them as a user preference profile; (4) receiving a query from the user at the third-party website associated with the market of the third-party; and (5) providing an advertisement to the user, wherein the advertisement is based on the ascertained preferences of the user. In embodiments, the determining preferences may be through the use of natural language processing. The advertisement may be provided by the computer facility. The advertisement may be provided through the third-party and enabled through preferences provided to the third-party from the computer facility. The API may enable the collection of at least one of cost information, product information, personal information, and topical information. The decision may be also based on information inferred from a user's social network. An advertisement may be delivered to other users associated with the user, such as through a social network.

[0189] In embodiments, the present invention may provide a taste and preference API that third parties may use to provide users with reviews from similar users, where the user and similar users may not have engaged in a dialog or interacting directly with the computing facility of the present invention. For instance, the taste and preference API may enable the present invention to collect taste and preference information for the user, provide the third-party with

taste and preference information for the user from a previously established taste and preference profile, provide the third-party with taste and preference information for the user based on near-term actions of the user, and the like. In embodiments, the user may have never interacted with a facility of the present invention, where a user's taste and preference profile may be created and updated through the user's interactions, responses, recommendations, reviews, and the like. In embodiments, the system may learn about a known user's taste profile through their rating things they like and don't like, or through the use of natural language processing, such as inferring a taste profile by analyzing how the user tags their user profile. In this case, similar users may have a previously established taste and preference profile, and as such may have a taste and preference profile that may be matched to the user. In addition, these similar users may have reviews associated with their profile. The system may now match the user to a similar user, and then provide the user with the associated review. For instance, a user may have an existing taste and preference profile with the system, such as directly with the computer facility or through at least one third-party API, and may want to know what other similar users thought of some product, service, person, event, and the like. The system may then search the taste and preference profiles for similar users on the subject the present user has interest in. In this way, the system may now be able to provide reviews and such to the present user from similar users, and thus helping the present user determine what they may want to do based on their tastes and preferences. For example, a user may go to a product website that utilizes the taste and preference API of the present invention, and is interested in reviews for digital cameras. The third-party may now find similar users, and then search for digital camera reviews by those similar users and provide the reviews to the current user. In embodiments, the reviews may be resident at a third-party facility, at another third-party facility, at a facility of the present invention, and the like. In embodiments, the ability to show reviews of similar users may allow the user to access more relevant reviews in a more time efficient manner, and the third-party user of the API may be able to provide more targeted and relevant support to their users.

[0190] In embodiments, the present invention may provide a taste and preference API that third parties may use to provide users with reviews from

similar users, where the users are determined to be similar without the similar users participating in a dialog through the present invention. For instance, the similar user may be identified as being similar through a social network, friend, family, work, and the like. In an example, a user may be associated with a second user through a social network, and through this association, determined to be 'similar', such as through age, interests, and the like. The similar user may then provide a review, such as for a product, an activity, and the like. This review may then be provided to the user as relevant through the similarity to the other user. In embodiments, similar users may be determined through similar recommendations on other topics, such as in combination with other factors, where the other factors may be a social association.

[0191] In embodiments, the present invention may provide for a computer program product embodied in a computer readable medium that, when executing on one or more computers, helps a user find reviews of similar users through the use of a computer facility by performing the steps of: (1) providing a user preference learning API to a third-party website to determine preferences of the user, wherein the preference learning API is executing as an extension of the computer facility; (2) collecting preferences of a plurality of users, wherein the plurality of users includes the user, (3) storing the preferences of the user in a taste and preference database which contains a plurality of taste and preference profiles; (4) receiving a request from the user through a third-party taste and preference learning API for a topical review from a user who has similar taste and preferences; (5) matching the preferences of the user to at least one other user's preference in the taste and preference database; (6) searching for a review related to the request for the topical review from amongst the matched other users; and (7) providing the review to the user. In embodiments, the determining preferences may be through the use of natural language processing. The review may be found within the computer facility, a facility of the third-party, and the like. The computing facility may be a machine learning facility. The preference learning API may enable the collection of cost information, product information, personal information, topical information and the like. The review may be provided by a similar user that has no taste and preference profile, where the user may be similar as determined through a social association, where the social association may be a social network.

[0192] In embodiments, the present invention may provide a taste and preference API that third parties may use to recommend products, services, and the like. For instance, a user may come to a third-party website in search of a recommendation for a product, and the third-party may then utilize the taste and preference API to better understand what the user typically prefers, and from that preference, suggest a product. In an example, the user may come to an audio store website looking for a recommendation for an audio system for playing music from their iPhone. The third-party may then utilize existing taste and preferences for the user through the API. In this example, the tastes and preferences of the user may indicate that they are a collage student and often on the run in their social life. From this information, the third-party may now make recommendations, such as recommendations for audio systems that are portable, small, powerful, and the like. Alternately, the third-party may use the taste and preference API to determine their taste and preferences at the time of the user inquiry, such as targeted to the inquiry, content of the third-party, for the user in general, and the like. The third-party may use this new taste and preference information alone, or in combination with previous tastes and preference profiles through the present invention, to make the recommendations. The taste and preferences as established through the third-party may now be stored in a facility of the present invention, such as to be used again or in combination with new taste and preference profiles generated through other third-party APIs or directly through a facility of the present invention. In embodiments, the ability to use the taste and preference API may improve the recommendations for products, services, and the like that are made through third-party sites.

[0193] In embodiments, the present invention may provide a taste and preference API that third parties may use to recommend products, services, and the like, to a user based on the actions of similar users. For instance, two users may have previously established taste and preference profiles with the present invention, where one of the users has selected a product, service, or the like, and where third-party may now provide a recommendation to the other user based on their similarity, such as determined through their profiles. In an example, two users may have been determined to be similar through their taste and preference profiles, such as by their age, location, political views, social activities, and the like. The first user may then select a product, such as a car. In the event that the

second similar user should indicate an interest in a car, such as through searching, advertisement selection, explicit question to the present invention, and the like, the present invention may provide the car selection to the second user as a potential fit because of their similarities. In embodiments, this may be done without a dialog provided to one or both users.

[0194] In embodiments, the present invention may provide for a computer program product embodied in a computer readable medium that, when executing on one or more computers, helps a user find a recommendation through the use of a computer facility by performing the steps of: (1) providing a user preference learning API to a third-party website to determine preferences of the user as applied to the products and services of the third-party, wherein the preference learning API is executing as an extension of the computer facility; (2) receiving third-party information related to the products and services of the third-party; (3) collecting the preferences of the user and storing them as a user preference profile, wherein the source of the collecting is from user interactions on the Internet; (4) receiving a query from the user at the third-party website associated with at least one of products and services of the third-party; and (5) providing a recommendation for at least one of the product and service to the user from the computing facility, wherein the recommendation is based on the query and the ascertained preferences of the user. In embodiments, the determining of preferences may be through the use of natural language processing. The collecting may be from the third-party website on the Internet; a plurality of third-party websites on the Internet; at least one of recommendations, purchases, and search result choices made by the user; and the like. The computing facility may be a machine learning facility. The third-party information may consist of product information from product manufacturers, product information from web merchants, pricing information from other websites, availability information from other websites, pricing information from merchants, availability information from merchants, a review, comments, ratings, and the like. The API may enable the collection of cost information, product information, personal information, topical information, and the like. The preferences may be derived from an action of a second similar user, where the similarity may be determined through a taste and preference profile for the user

and second similar user. The action of a second similar user may be a selection of at least one of product and service.

[0195] In embodiments, the present invention may provide a taste and preference API that third-party social network sites may use to show a user people that are similar to them on the social network. These similar people may be shown as a list, as photos, by region, by age, by gender, and the like. For instance, the user may come to a social network site and ask to see or be connected to people who are similar to them. The social network site may then utilize the taste and preference API to provide the user with a dialog to determine their tastes and preferences, such as in general, to social situations, to social networking, to activities, to music, to personality, and the like. Alternately, the user may already have a taste and preferences profile as determined directly by a facility of the present invention, through another third-party API, through the social network site, and the like. The social network may then use this information to match the user to other people on the social network, such as through taste and preference profiles of the other people as previously determined, through information available about the other people as available through the social network, and the like. For example, the user's taste and preference may indicate that they are young and enjoy going to clubs in the NYC. The social networking site may now be able to match the user to similar people on the social network, such as by list, photograph, by category, by region of the city, and the like. In embodiments, the taste and preference API with a social network may provide an enhanced matching experience to the user who is trying to find other similar people to be social with.

[0196] In embodiments, the present invention may provide for a computer program product embodied in a computer readable medium that, when executing on one or more computers, helps a user find other similar users on a social networking site through the use of a computer facility by performing the steps of: (1) receiving an initial request from a user through a third-party social networking site API, wherein the initial request is to find other similar users to them on the social network; (2) ascertaining preferences of the user through the social networking site API; (3) matching the user preferences to other users on the social network with users with similar preferences; and (4) providing a matching result to the user that includes the other users that match user's

preference. In embodiments, the ascertaining of preferences may be through the use of natural language processing. The matching result may be presented to the user as a list of the similar users. The matching result may be presented to the user as a profile of the similar users. The matching result may be presented to the user as links to the similar users within the social network. The computing facility may be a machine learning facility.

[0197] In embodiments, the present invention may provide a taste and preference API that third-party searching facilities may use to rank search results based on which results similar users selected the most. For instance, the searching facility may offer users the opportunity to improve the relevancy of how the search results are listed through a tastes and preference profile, as provided through the taste and preference API of the present invention. A taste and preference profile database or the like may then be accumulated and maintained, from which the searching facility may rank search results for the user to previously selected results by other similar users. In an example, a user may have a taste and preference profile that shows a retired male who likes to sail and is a bit adventurous. When the user searches for Caribbean vacation destinations the searching facility may rank the search results with these taste and preference attributes listed first, such as for sailboat rental packages in the islands, hiking in the islands, off-beat destinations, and the like. In embodiments, the use of a taste and preference API provided to a searching facility may improve the relevancy of ranked search results to the user.

[0198] In embodiments, the present invention may provide for a computer program product embodied in a computer readable medium that, when executing on one or more computers, helps rank search results through the use of a computer facility by performing the steps of: (1) receiving a search request from a user through a third-party searching facility; (2) ascertaining preferences of the user, wherein the ascertained preferences from the user creates a taste and preference profile for the user and is stored in a taste and preference storage facility which includes a plurality of other user taste and preference profiles, where the profiles also contain a history of search results selected by the other users in previous searches; (3) matching the user to other users with similar taste and preference profiles; (4) determining a search result set for the user's search request; (5) matching the search result set to the history of search results

selected by the other users with similar taste and preference profiles; and (6) providing the search results to the user, wherein the search results are ranked according to the matched results selected by the other users with similar taste and preference profiles. In embodiments, the ascertaining of preferences may be through the use of natural language processing. The computing facility may be a machine learning facility. The searching facility may be a search engine.

[0199] Referring to Fig. 34, the present invention may utilize social graphs to infer the taste and preferences for an unknown user by finding paths through an Internet based social interactive construct to people with known taste preferences. In this way, the present invention may provide a way to get data for a user that the system has never heard of before. In embodiments, the present invention may provide for a computer program product embodied in a computer readable medium that, when executing on one or more computers, helps determine an unknown user's preferences through the use of internet based social interactive graphical representations on a computer facility 3402 by performing the steps of: (1) ascertaining preferences of a plurality of users who are part of an internet based social interactive construct, wherein the plurality of users become a plurality of known users 3404; (2) determining the internet based social interactive graphical representation 3412 for the plurality of known users 3408; and (3) inferring the preferences of an unknown user present in the internet based social interactive graphical representation 3412 of the plurality of known users based on the interrelationships between the unknown user and the plurality of known users within the graphical representation 3410. In embodiments, the Internet based social interactive graphical representation may be a social network, a social graph, a social diagram, and the like. The unknown user may be three degrees, five degrees, and the like away from the closest known user in the internet based social interactive graphical representation. The inferred preferences of the unknown user may make the unknown user a new known user, and the new known user may be used to contribute to the inferring of the preferences of a second unknown user. The preferences may include personal information, topical information, and the like related to interactions of the user, where the interactions may be through the internet based social interactive graphical representation. The interactions may be through an API provided to a third-party website. The inferring may be provided in conjunction

with other known users that are related to the user in the user's internet based social interactive construct. The ascertaining of preferences may be through the use of natural language processing. The computing facility may be a machine learning facility. The inferred preferences may be used to target advertising to the unknown user, to share reviews with the unknown user. The inferred preferences may be used to recommend products, services, and the like to the unknown user. The inferred preferences may be used to aid in ranking search results for the unknown user. Known users that are in close proximity to the unknown user may carry more weight in an inferring algorithm. The inferred preferences may be refined by information from other sources, where the other sources may include third party sources, recommendations made by the plurality of known users, search queries by the plurality of known users, search result selections one of the plurality of known users, personal tastes as determined through web interactions by at least one of the plurality of known users, and the like. The other sources may include a third-party preference learning API.

[0200] In embodiments, the present invention may utilize social network graphs, diagrams, graphical representations, and the like, to infer the taste and preferences for an unknown user by finding paths through a social network to people with known taste, or visa versa. Social diagrams are, generally speaking, the mapping of a plurality of users and how they are related. By using the social diagram, a taste and preference of a known and unknown user may be determined from their interrelation within the diagram. For instance, a user with a known taste and preference profile may be directly associated with a plurality of other users, such as represented in a social diagram. To a first approximation, it may be assumed that these plurality of other users are similar to the user, and so have similar taste and preferences. These other users may then be provided refined services that take advantage of knowing a users' taste and preference, such as described herein. For example, provided with a user with a known taste and preference profile indicating they are a rock climber, it may be assumed that users within a first link of the user's social diagram are also rock climbers. In reality, this may prove to be too general an assumption. However, it may be a good assumption that the user does have associations with other rock climbers, and so the system may go out through the social diagram searching for other known users that enjoy rock climbing. In this example, it may be found there is

another known user, such as three links away, that also enjoys rock climbing, and this user is found in a cluster that connects to the first user. From this it may be a good assumption that this cluster is a group of rock climbers, and rock climbers may all share a set of similar tastes and preferences with each other. In embodiments, tastes and preferences may be inferred from associations within a social network diagram, and as such, may be provided benefits from the present invention as described herein.

[0201] In embodiments, the present invention may provide for a computer program product embodied in a computer readable medium that, when executing on one or more computers, helps determine an unknown user's tastes and preference through the use of social network graphical representations on a computer facility by performing the steps of: (1) ascertaining preferences of the user, wherein the user becomes a known user; (2) determine the social network graphical representation for the known user; (3) determine the presence to other known users within the known user's social network graphical representation; and (4) infer the preferences of an unknown user present in the known user's social network graphical representation based on the interrelationships between the unknown user and the known user and other known users within the network graphical representation. In embodiments, the ascertaining of preferences may be through the use of natural language processing. The social network graphical representation may be a social graph, a social diagram, and the like. The computing facility may be a machine learning facility.

[0202] In embodiments, the present invention may combine the tastes and preferences of a user as determined through two or more third-party API to improve recommendations provided through the two or more third-party API. For instance, there may be taste and preference profiles being established through more than a single third-party API, and by combining these different tastes and preference profiles by the present invention, a combined taste and preference profile may be generated. Further, as additional taste and preference profiles are created through third-party API, they may be used to continuously update the combined taste and preference profile for a user. The third parties may then utilize the combined taste and preference profile to improve their recommendations. This may especially be the case when different third parties focus taste and preference profiling on different areas, such as products, personal

relationships, services, celebrities, and the like. It can be appreciated that combining a number of more specific profiles into a combined profile may provide a richer taste and preference profile than could be generated through any one of the more specific profiles. In addition, a user may change their tastes and preferences over time, and so combining more recent user profile interactions on one third-party API may benefit another third-party that the user has not interacted with in recent time, but where the other third-party wants to keep their user profiles up to date.

[0203] Referring to Fig. 35, in embodiments the present invention may provide for a computer program product embodied in a computer readable medium that, when executing on one or more computers, provides improved taste and preference profiling through the use of a computer facility 3502 by performing the steps of: (1) creating a first taste and preference profile of a user 3512 through the user's interactions with a first third-party website 3514 through a first preference learning third-party API 3504; (2) collecting additional user interaction information through a second third-party website 3514 through a second third-party API 3508; and (3) combining the additional user interaction information with the taste and preference profile to improve the taste and preference profile 3510. In embodiments, creating a first taste and preference profile may be through the ascertaining of user preferences through the use of natural language processing. The computing facility may be a machine learning facility. The API may enable the collection of at least one of cost information, product information, personal information, and topical information.

[0204] Referring to Fig. 36, in embodiments, graph constructs 3620 may be developed and/or utilized by a web-based advice facility 3602, such as to aid in providing recommendations to users 3608 through a dialog with the user 3608 across the Internet 3604, with a minimized amount of dialog with a user, to provide recommendations to users where the graph constructs augment the process that leads to the recommendation, where the use of the graph construct eliminates the need to carry on a dialog with the user to form recommendations to users, and the like. In embodiments, graph constructs may be developed through information from third-party sites 3610. In embodiments, there may be a plurality of types of nodes 3612 in the graph, such as people, entities, tags, and the like. For instance, people may be users of websites, applications, mobile

devices, shoppers in a store, anonymous web browsers represented purely by a unique cookie id, and the like. Entities may be things that people like, dislike, buy, search for, research, and the like. Tags may be short textual descriptions of entities, people, and the like. In embodiments, nodes in the graph may be connected by a plurality of types of edges 3614, such as for preference data, tagging data, and the like. For instance, entities may be connected to people nodes by preference edges that express the degree to which a person likes or dislikes that entity. Tags may be connected to people and entity nodes by whether those people or entities are tagged with those tags.

[0205] In embodiments, graph data may be explicitly given by users (e.g. user 'A' says they like thing 'B'), crawled from publicly available web sites, provided by third-party sources, and the like. Once data is received, the system may attempt to "alias" it to existing data in the system. For example, if the data tells the system that user 'A' likes restaurant 'B', then the system attempts to identify what, if anything, is already known about restaurant 'B' through things like matching names, addresses, phone numbers, and other information. This may allow the system to aggregate data, such as training data, from multiple sources all against the same entity representing restaurant 'B'. The system may perform aliasing against users. For example, user jsmith99 might be the same user as john_smith on two different websites. The system may use similarity of usernames, email addresses, pictures, full first and last names, geographic location, and the like to correlate users across different web sites and identity systems.

[0206] In embodiments, nodes in the graph may have a "taste profile", such as with a numerical quantity. A person may be predicted to like or dislike an entity or tag based on their taste profiles. Similarly, two people may be predicted to be similar or dissimilar based on their taste profiles. Entities may also be compared to see how similar they are to each other using their taste profiles. The system's graph may initially have taste profiles assigned to some nodes and then propagate those taste profiles to the nodes that don't have taste profiles. This propagation may be an iterative process that "flows" taste profiles from nodes that have profiles into nodes that do not have profiles. Alternatively, the iteration may update the profile of nodes that already have a profile based on neighboring nodes' profiles. New data may be incorporated into the graph by

adding new nodes or edges and then updating the new or changed node purely using neighboring nodes' profiles. Alternatively, the system may run multiple iterations of updates across the entire graph.

[0207] Many different kinds of data may be fit into being viewed as a "like" or "dislike". For example, viewing a web page can be represented in the graph as an edge with a weak connection between the person viewing the web page and an entity representing the web page. Someone buying a book can be represented by an edge making a strong connection between the person buying the book and the book itself. Someone answering a question that has three mutually exclusive answers can be represented as an edge between the person and an entity representing the answer they gave as well as two negative edges to the two answers the user did not give.

[0208] In embodiments, methods and systems may provide for recommendations to users based on the degree to which a recommendation may be new, interesting, and the like, which will herein be referred to as 'interestingness'. In embodiments, interestingness may be a combination of being an interesting subject, topic, product, and the like, as well as how new or revived the idea is. In an illustrative example, the user may live in the U.S. and be interested in cooking Italian food, and so the system may provide cooking recommendations to the user. In this instance, recommending a 'new' cheese flavor to try as Parmesan Cheese may have a low interestingness, because the use of Parmesan Cheese in the U.S. may not be new at all, and even be rather over used. Alternatively, Pecorino cheese is an Italian cheese made from sheep's milk, and can be used instead of Parmesan cheese on pasta dishes and is sometimes preferable if a sharper taste is desired. As such, a recommendation to substitute Parmesan with Pecorino may be considered to have a high interestingness, at least in a relative sense to that of Parmesan. Interestingness may be determined relative to what is standard or typical, relative to a past recommendation, relative to a novel factor, and the like. In this instance, the interestingness of Pecorino may be rated high relative to the interestingness of Parmesan.

[0209] In embodiments, interestingness may be correlated to a known taste profile and fresh to the user, where fresh may mean new (such as new to the world), absent from a user's own past experience (such as by reference to a

known history of the user), and the like. Alternately, 'fresh' may not necessarily be new to the world or to the user, but be a new fact or story associated with something that makes it interesting anew. For example, "Chipotle on 21st street" may not be new or interesting, but if someone provides a recommendation to "Get a burrito at Chipotle where Oprah gets her burritos" then it's interesting. The system may encourage this by requiring users to write a reason as to why they are recommending something. Interestingness may be determined as related to the ratio of users rating the item, to the item being "saved" by users. For instance, if there is a bookmark, save for later, add to wish list, and the like functionality (e.g. a star rating functionality), it may be seen that there is a correlation between highly rated items that are 'low saved' items as not being interesting. In this instance, it could be that everyone knows about the items, so they can rate them, but they're not worth saving for later, thus not interesting. Alternately, an item that is highly rated and often saved by users may be considered interesting, because the item is both highly regarded (i.e. highly rated) and worth saving for further consideration. Interestingness may be an acceleration of social activity. For example, there may be a restaurant that has been around forever, and that the user knows about, but which on a sunny weekend suddenly starts getting a lot more foursquare.com check-ins. This may be a sign there's something interesting going on there now. In embodiments, the advice facility may also determine that something is new or new to the user by looking at release dates of books, movies, albums, products, and the like, and take the earliest date the item is found on the Internet to determine its interestingness; look at the date of the first review written for something on the web; look at events such as movie releases, concerts, author talks, and the like, that may be considered inherently new; and the like. The advice facility may also ask the user to rate things they already know about, where the system may assume that items user's haven't rated are new to them. The advice facility may have a 'save' feature to encourage users to use when they don't know about something yet but want to check it out.

[0210] Thus it will be understood that 'interestingness' as that term is used herein may include (or more concretely, be quantitatively evaluated according to) relevance in the conventional sense, particularly as it relates to the relationship between a user's profile (or taste profile) and new content. A wide

variety of analytical, mathematical, rule-based, and/or heuristic techniques are known for evaluating relevance, any of which may be usefully adapted to determining relevance (and more generally interestingness) as contemplated herein. However, interestingness additionally includes dynamic relationships between a user and content based on, e.g., time, location, user history, and so on.

[0211] Time, for example, may be important simply as a measure of newness, such as where a current statistic is more interesting than an older measurement of the same statistic. Conversely, where a user expresses interest in a particular point in time or period in history, older statistics, facts, opinions, and the like having an explicit time (as determined by metadata, content, chronology, or the like) may be more interesting. Newness - that is, a measure of how recent an item is - may be particularly important to interestingness where there are numerous diverging items of information on a topic and there is a measurable increase in the current popularity of or interest in particular ones of the diverging items. This type of popularity may be measured in numerous ways such as passive measurements of blogging activity, newly indexed web content, or any other Internet-based measurement of user interest, as well as active measurements of hits, traffic, or other activity at web servers, as well as group or individual monitoring of client activity. Time may also be important to interestingness of an item in other contexts, such as where time is explicit or implicit in a user inquiry, e.g., things to do this weekend, movies showing this evening, etc.

[0212] Location may also significantly impact interestingness. This may include simple geographic proximity using any suitable location-aware technologies, and may incorporate other aspects of a user profile such as an interest in particular venues (e.g., food, art, entertainment) or a current activity associated with a user. However, it will be understood that this may also include location-related items such as inferences about the convenience of adjacent locations through various transportation alternatives available to the user (e.g., a car, public transportation, etc.), as well as a user's available budget for immediate or extended travel planning. In addition, the context of a location and its corresponding interestingness may depend on other dynamic location attributes such as the location of friends within a social network, and the proximity to or distance from geographic concentrations of the same.

[0213] User history may also be used to parameterize interestingness. For example, where a new item is responsive to a user inquiry or well matched to a user profile, but highly distinct from previous content obtained by a user, this distinctiveness may make the item more quantitatively interesting even if the calculated relevance is equal to or less than relevance of other results. Thus in one aspect, interestingness may depend concurrently on measures of similarity (or relevance, or the like) and dissimilarity, or more specifically, characteristics that make an item dissimilar to previous content in a user's history. Alternatively, an item of information may rank poorly on a general measure of relevance that is de-emphasized based on other aspects of the user's current context. Thus interestingness may provide a measure of relevance to a user based on any suitable similarity or matching metric that is further augmented by a newness to a user, as explicitly measured through a dissimilarity to information in a user's history, or any objective basis for adapting relevance scoring based on the user's context. In one aspect, interestingness may be objectively measured as relevance based on a user's profile along with dissimilarity to a user's history and one or more aspects of a user's current context such as time or location. As measured in this manner, many objectively highly relevant items may not be particularly interesting to a user, while marginally relevant items may be highly interesting.

[0214] Referring to Fig. 37, recommendations may be provided through a recommendation facility 3704 as part of a web-based advice facility 3702. In embodiments, the recommendation facility may utilize interestingness filtering 3708 in the process of generating recommendations to a user 3722. Recommendations sources may include friends 3718, similar users 3714, influential people 3720, source websites 3712, and the like. Recommendations may be provided to the user through the Internet 3710, through a telecommunications facility 3724 (e.g. cell phone network), and the like.

[0215] In embodiments, the determination of interestingness may be related to 'social activity' of other individuals (e.g. friends, famous people, an authoritative person), the 'born on' date of a product, place, event (e.g. the opening of a restaurant, the release of a movie, a new product), and the like. The social activity of other individuals may be related to individuals influential to the user, such as friends highly rating a topic, friends with similar tastes highly rating a topic, non-friends who are influential highly rating a topic, non-friends

who have similar tastes in this topic highly rating it, and the like. These other individuals may fall into different categories, such as friends; people the user doesn't necessarily know but whom have similar tastes in this topic; people who don't necessarily have the same tastes as the user but who are famous, prolific, well-known; critics in this topic (e.g. movie reviewers); and the like. The system may also note which of the user's friends have similar tastes to them in this specific topic. For example, if Ted has similar restaurant tastes to the user, but Alice does not, then a restaurant may be interesting to the user if Ted likes it but not necessarily interesting if Alice likes it. Other reasons that something may be considered interesting for the user are if a friend has saved the recommendation for later, if friends are commenting and discussing the recommendation actively, and the like.

[0216] In embodiments, recommendations may be provided to users through email, social networks, third-party sites, when the user requests, as a data feed, as a push service, on a periodic basis, in association with a search topic, related to a current geographic location, to a home computer, to a mobile computer, to a mobile communications facility (e.g. cell phone, smart phone, PDA), and the like. For example, the user may be provided interestingness recommendations to their mobile phone based on their current geographic location, such as product recommendations to stores in the area (e.g. products on sale, new products, products that are difficult to get), places to see, restaurants to try, and the like, where the recommendations are based on interestingness. In this way, the user doesn't just receive recommendations, but rather a more interesting set of recommendations, which may increase the chances that the user will be interested in the recommendation, such as in a 'discovery' of a new idea, place, product, and the like. And when the discovery is associated with a particular genre, becomes a 'local discovery', a 'restaurant discovery', a 'technology discovery', a 'cooking discovery', and the like.

[0217] In embodiments, the present invention may provide 'local discovery' to a user, where local discovery may include providing new and interesting things to the user instead of relying on the user typing something into a search box or otherwise "pulling" search results to them. This may be especially useful for mobile devices where typing is more difficult, such as when a device input is input constrained (e.g. small keyboard, small display, the user

being mobile (walking, driving), and the like). Although the description of local discovery herein is provided primarily with respect to a mobile device application, one skilled in the art will appreciate that it may be implemented on any computing facility, such as a laptop, desktop, navigation device, or the like. Local discovery functionality may also be available through a web interface, through emails of "new stuff" for the user, through Twitter, through posting to blogging platforms (e.g. Wordpress, Tumblr, etc), and the like. Further, local discovery content may be provided to a user upon request, transmitted to the user (e.g. email) to push new interesting things to a user (e.g. weekly), and the like.

[0218] In embodiments, a mobile device local discovery application may show the user places nearby, such as that they've rated in the application in the past, that their friends have rated, that people with similar taste as the user have rated, that authoritative sources have rated, that famous people have rated, and the like, where 'rating' may be an inferred recommendation from behaviors (e.g. online or off-line) of the person. In embodiments, the places that get shown may be restaurants, bars, boutiques, hotels, and the like. There may also be a navigational element to let the user filter down to narrower lists, such as for example, "Italian Restaurants" that are nearby or at a specified location, such as recommended by others.

[0219] Besides showing places, local discovery may also show the user items to buy, events to go to, things to look at (such as if they have some time to kill), and the like. For example, local discovery might provide a list of recommended books the user might want to read and optionally show where to buy them locally. Local discovery might select the items, provide recommendations, and the like, based on machine learning, such as what the user's friends have liked recently, what people with similar tastes in books have liked recently, what popular/prominent critics have liked recently, and the like, or just what is popular overall or popular near the user. Similarly, this also applies to other kinds of products, events to go to, and the like.

[0220] In embodiments, local discovery may find people with similar tastes in each area (restaurants, books, etc) and then let the user follow the things they rate. When using the mobile application, local discovery may use the user's location to filter down lists of things people similar to the user like near by. Local discovery may determine whether another person has similar taste as the

user through machine learning, asking both to rate various places and things, asking both to answer questions to gauge similarity, and the like. Local discovery, such as through the advice facility, may then try to validate similarity between the user and another person, such as based on liking obscure things in common, disliking popular things in common, showing written reviews that the other person has written, describing the other person's traits (demographics, location, etc), showing how many other people follow the other user, and the like.

[0221] In embodiments, instead of producing a list of recommended places, things, or events for a given area, local discovery may also produce a "discovery" feed of interesting stuff for the user, such as with a high interestingness rating. This may mean that instead of seeing the same ten restaurants recommended every time the user looks near their office, they may see a few different results each day show up. Ideally these new restaurants may be showing up based on the user's friends or people with similar tastes liking some new place near you, but it may also be a partially editorial process where staff members associated with the implementation of a local discovery application are constantly finding new places and sending them out to the user base.

[0222] In embodiments, the user may also save things to a "wish list", "to do list", and the like, for using later. This saved list of products, places, events, and the like, may then also be used to alert the user about deals, availability, new reviews, and the like, about those products, places and events so as to only alert the user about stuff they're interested in. For example, it might be annoying if the user's phone vibrated to tell them some shoes were available at a store they were walking by unless they had previously indicated they wanted those shoes on their "wishlist".

[0223] In embodiments, the user interface for a local discovery application may be a map, a textual list, a "cover flow" like interface for flipping through (such as in Apple Computer's implementation of cover flow), and the like. The interface may also send alerts to the user when a friend or someone with similar taste likes something nearby, likes something that the user has saved to their wish list, and the like.

[0224] In embodiments, a local discovery application may be associated with the weather, where the application may in part determine recommendations based on actual or predicted weather in the area the person wants a recommendation in. For example, the application may recommend places with great outdoor seating when the weather is warm and not raining, recommend things to do outdoors when the weather is good, provide more weight to recommendations for going to a museum when the weather is bad, and the like.

[0225] In embodiments, a local discovery application may be associated with the time of day, such as taking the hours of operation and the distance to the place into account when making recommendations. For example, if the person wants a place to eat now, the system may not recommend a place that is closed or about to close. Similarly, the system wouldn't recommend things to do that have sold out already, that will not be reachable in time using estimated transportation time to reach them, and the like.

[0226] In embodiments, a local discovery application may be associated with taste, such as using a person's taste when making recommendations instead of just showing what's popular among users nearby. In embodiments, taste may be inferred through the things the user likes, the people they follow on social networks, and the like, such as described herein.

[0227] In embodiments, a local discovery application may be associated with location, such as using the person's location when making recommendations, use their location when offering the user a list of topics they can get a recommendation in, and the like. For example, if there are no hair salons near by, the system would not offer the user the choice of getting recommendations about hair salons. Similarly, if there are no Mexican restaurants or video game stores, the system would not offer the choice of recommendations in those topics. If the user's location is inside a store, the system wouldn't offer recommendations about topics the store is out of or does not sell.

[0228] In embodiments, a local discovery application may be associated with social activity, such as showing recommendations that have received some amount of activity from friends, from respected authorities in the topic, from people with similar tastes as the user, and the like.

[0229] In embodiments, a local discovery application may be associated with recommendations that are interesting in that they are new, newly popular, have received social activity, have an indication that they are not new but also not popular enough that the user is likely to already know about them, and the like, such as described herein in terms of interestingness.

[0230] In embodiments, a local discovery application may be associated with what items a user saves on their mobile device, such as when a person using a mobile application may "save for later" interesting recommendations as well as things they see in stores or while out walking around. For example, if they see a book they like in a book store, they may scan the bar code and save it for later. If they see a restaurant they like they may take a picture and capture their location, such as via GPS, and save it for later. Saved content may then be the basis of recommendations used in a mobile application later on, or the content may be sent to the user through other channels, such as a weekly email reminder, through a web application, and the like. Saving may also be used by the system as a social indicator, such as to indicate that something is interesting content for other users.

[0231] In embodiments, a local discovery application may be associated with what items a user saves on the Internet. This may be similar to saving on a mobile device itself, where the user may save content they find on the web and then use it later through their mobile application. For instance, if the user sees a book review and saves it for later they may then be reminded about it later when they use their mobile device to look for book recommendations. Similarly they may save a restaurant or thing to do and later be reminded about it on their mobile device.

[0232] In embodiments, a local discovery application may be associated with displaying topics to offer recommendations in, such as displaying a grid of pictures when the mobile application first starts representing the topics that recommendations are available in. The choice may be based on the user's location, their historic use of the application, and the like, where the application predicts what the user is interested in. For example, there might be pictures of restaurants, iPhone apps and video games among others when the application starts based on their being recommended restaurants near the user and the application's belief that the user is interested in iPhone apps and video games.

The actual pictures may be picked based on the application's knowledge of the user. For example, the picture for the restaurants topic may be a picture of a restaurant nearby that the application predicts the user may like. The size and sorting of the pictures may be based on how interested the application predicts the user will be in that topic so that the user's most frequently used topics are at the top of the list and shown with the biggest picture, such as shown in Fig. 38. The user may then select one of the recommendations, and be linked to more detailed information associated with the recommendation, such as shown in Fig. 39.

[0233] In embodiments, recommendations may be displayed as the images with an ordering as a function of interest to the user, such as in an irregular grid where the left-right top-to-bottom ordering is based on how much the system thinks the user will be interested in each recommendation or each topic. For example, as shown in Fig. 40 the system has predicted that the user is more interested in restaurants than coffee shops and so the restaurants tile is listed before the coffee shop tile. Also, the content of each tile is a recommendation in that topic that the user may like. So the restaurant tile shows a restaurant the user may like that is near them, the movie tile shows a movie the music may like, etc.

[0234] In embodiments, a local discovery application may be associated with determining a radius to get recommendations from, such as the application has to pick a radius of how far from the user's location recommendations will be returned in. The radius may be selected by the user, based on the population density of the area around the user, and the like. For example, in NYC the radius might be .025 miles while in rural SC the radius might be 60 miles.

[0235] In embodiments, a local discovery application may be associated with determining which people to show reviews from, such as when recommendations are shown in the application they may be accompanied by reviews, ratings or other recommendations from people. The application may choose which people to show based on whether they have similar taste as the person using the application, whether they are friends of the user, whether they are authoritative critics, and the like, such as shown in Fig. 41.

[0236] Referring to Fig. 42, methods and systems may provide a recommendation to a user through a computer-based advice facility 4202,

comprising collecting topical information, wherein the collected topical information includes an aspect related to the extent to which a topic is interesting, or an interestingness aspect; filtering the collected topical information based on the interestingness aspect 4204; determining an interestingness rating from the collected topical information, wherein the determining is through the computer-based advice facility 4208; and providing a user with the recommendation related to the topical information based on the interestingness rating 4210. In embodiments, the interestingness aspect may be derived at least in part from social activity of another individual that indicates a recommendation for a topic. The other individual may be a friend, a famous person, an authoritative person, or and the like. The other individual may have similar tastes to the user, either in general, or with respect to a particular category or type of interest. The social activity may be saving a recommendation. The social activity may be at least one of commenting and discussing a recommendation actively. The social activity may be collected from activity of the other individual with respect to sources on the Internet, such as social networking activity. The interestingness aspect of the topical information may also be determined based upon an indication that the topical information has some aspect of freshness or newness. The newness may be an indication that the topical information is new topical information over a predetermined period of time. The newness may be an indication that the topical information is newly popular. Status as newly popular may be determined from an activity level on the web. Newness may be general, such as the emergence of a newly popular topic within a social network, or it may be particular to a user, such as when an older topic is first exposed to a user, making it new to the user, if not to the social network as a whole. The interestingness aspect of the topical information may be based at least in part from at least one of a review, a recommendation, a blog entry, a tweet, an authoritative source, a news source, an e-publication, a purchase, a view, a time viewed or and the like. Without limitation, the interestingness aspect may be based on time data. The time data may be a release date, such as a movie, a product, and the like. The time data may be an event opening, such as a restaurant opening, a cultural event opening, and the like. The interestingness aspect may be frequency data, such as relating to how frequently the topical information is referenced in online sources. The interestingness

aspect may be related to a user interaction with a computer device. The user interactions may be interpreted by a machine learning facility as user behavior that indicates a preference level for the topical information by the user. The user interaction may be selection of a web link. The user interaction may be at least one of tapping, touching, and clicking on the computer device screen. The computer-based advice facility may include a machine-learning facility. The computer-based advice facility may include a recommendation facility. The filtering may be collaborative filtering. Recommendations may be sent to a user's mobile communications facility to provide recommendations in the user's current geographic area. There may be a graphical user interface on the mobile communications facility that provides the user with the ability to refine provided recommendations to the user. The recommendations may show the user at least one of items to buy, events to go to, things to see, and the like. The recommendation may be related to a local store. The recommendation may be related to a local restaurant. The recommendation may be related to a local bar. The recommendation may be related to entertainment. Recommendations may be further filtered to the user based on interestingness specific to the geographic area. A recommendation feed may be sent to the user for the current geographic area the user is located in. Recommendations may only be sent to the user that meet a threshold in confidence and in how much the system predicts the user will like the recommendation. The threshold in confidence may be related to the interestingness rating. The threshold in confidence may be determined by a machine learning facility based on past behavior of the user as related to previous recommendations provided by the system. The user may be able to save recommendations to storage on the mobile communications facility. The user may be able to save recommendations to storage with the computer-based advice facility.

[0237] Referring to Fig. 43, methods and systems may provide a geographically localized recommendation to a user through a computer-based advice facility 4302, comprising collecting a recommendation from an Internet source, wherein the recommendation is determined to have an interestingness aspect. The recommendation may further be determined based on a geographic location aspect 4304. Further options may include comparing the collected recommendation to a derived user taste and the user's current geographic

location 4308, determining at least one recommendation for the user based on processing on the comparison 4310, and delivering at least one recommendation to a user's mobile communications device, wherein the user is enabled to at least one of view, save, and share the recommendation, such as via an application at least in part resident on the computer-based advice facility 4312.

[0238] In embodiments, the computer-based advice facility may be a mobile communications device. The mobile communications device may be a smart-phone. The viewing may include providing source information from the Internet source. The source information may include the original recommendation, a rating, an image associated with the Internet source, or the like. The image may be a photo of an individual who provided the recommendation, or another indication, such as an icon, representative of such an individual. The source information may include a visual indicator of an extent of similarity to the tastes of at least one individual who made the recommendation. The indication may be qualitative (e.g., "this individual has tastes highly similar to yours") or quantitative, such as expressing a metric that measures relative similarity (e.g., "you share 10 interests out of 20 categories with this individual"). The source information may include a visual indicator of an extent of similarity to the tastes of more than individual who made a recommendation. The visual indicator may indicate a sorting of individuals based on the extent of taste similarity. Individuals may be listed in decreasing order of similarity.

[0239] In embodiments, the interestingness aspect may be topical information that is new, such as determined by a date of emergence of the information within a domain, such as an Internet domain, a collection of Internet news sources, an enterprise network, a social network, or the Internet as a whole.

[0240] The interestingness aspect may be topical information that is deemed by the advice facility to be new to the user, such as by comparison to past content reviewed by the user, accessed by the user, or the like, as reflected by a user's browsing history, by tracking the user's activities on one or more devices, or the like.

[0241] In embodiments the interestingness aspect may be topical information that has a new aspect to an existing topic, such as an update to a news item in which the user has shown interest in the past, such as reflected by

user feedback or by a user's activities, such as access to the item, time spent reviewing it, or the like.

[0242] The interestingness aspect may be determined as related to the ratio of users rating the item, or to the item being saved by users.

[0243] In embodiments the interestingness aspect may be determined by an acceleration of social activity associated with the topical information.

[0244] In embodiments the interestingness aspect may be determined based on information having more than one aspect, such as the information being determined to be similar to a user's profile (e.g., similar to items in which the user has shown past interest, matching a category of the user's interest, showing relevance or interest to other users who have similar tastes, or the like) while at the same time being dissimilar to a user's history (i.e., being new to this user in one of the ways noted above). Thus, for example, a user who has expressed a past in a particular celebrity might be expected to have very high interest in a breaking news item with respect to that celebrity.

[0245] In embodiments, the interestingness aspect may be further determined based on a user's current location and a temporal factor, where the temporal factor is based on topical information that is new as determined by one of the factors noted in this disclosure, such as being new to a domain (up to an including the entire Internet, but optionally being based on being new with respect to a domain that has a link to the user's current location), based on topical information that is deemed by the advice facility to be new to the user, based on topical information that has a new aspect to an existing topic, determined as related to the ratio of users rating the item, based on the item being saved by users, determined by an acceleration of social activity associated with the topical information, and the like. The geographical aspect may be a geographical location associated with the topical information, where the geographical location may be the location of an event, the location of a store, the location of restaurant, the location of point-of-interest, at least one product location, and the like.

[0246] In various embodiments the derived user taste may be based on a rating, where the rating may be provided by the user, friends of the user, people with similar taste as the user, an authoritative source, a famous person, inferred from user behavior, on machine learning with respect to user online behavior, and the like. The user behavior may be online behavior, including buying

behaviors, browsing behavior, social networking behavior, location-based behaviors, and the like. The recommendation may be items to buy, places to visit, events to attend, places to eat, and the like. The recommendation may be based on one of ratings and recommendations of at least one other user with similar tastes to the user from the current geographic location of the user. The similar tastes may be determined by machine learning through at least one of ratings from the other user and online behavior of the other user. The recommendation may be provided as part of a feed of local discovery recommendations. The recommendation may be saving of a recommendation saving to a list, where the list may be a wish list, to-do list, an events list, a deals list, and the like. A saved recommendation may be shown to the user through the local discovery application when the geographic location aspect of the saved recommendation matches the current location of the user. A recommendation may be forwarded to a user based on a recommendation from at least one other user, where the advice facility determines an applicability radius around the user's current location for use of the other user's recommendation. The application may be a local discovery application, where the local discovery application correlates at least one of new and saved recommendations with the weather, at least one of new and saved recommendations with the time of day, at least one of new and saved recommendations with the user's social activity, and the like. The local discovery application may display images based on how much the advice facility thinks the user will be interested in at least one of each recommendation and each topic. The displayed images may be displayed in an irregular grid where the left-right top-to-bottom ordering may be based on how much the advice facility thinks the user will be interested in at least one of each recommendation an each topic.

[0247] Turning now to Fig. 44, a recommendation system 4400 includes a matrix generation module 4402, a prediction module 4404, and a ratings database 4406. The modules of the recommendation system 4400 may be implemented on a single device such as a recommendation device, or on separate devices interconnected via a network. The modules of the recommendation system 4400 may be implemented on, for example, the machine learning facility 1318 illustrated in Fig. 13 and Fig. 15, the computing facility 1418 illustrated in Fig. 14, the machine learning engine 1510 illustrated in Fig. 15, the third party

website 3220 illustrated in Fig. 32, the third party website 3320 illustrated in Fig. 33, the first and second third party website 3514 illustrated in Fig. 35, web-based advice facility 3602 illustrated in Fig. 36, web-based advice facility 3702 illustrated in Fig. 37, and so forth.

[0248] According to various exemplary embodiments, the recommendation system 4400 may provide recommendations regarding different items to different users. The recommendation system 4400 may provide such recommendations by predicting a presently unknown affinity between a particular user and a particular item, based on some number of known affinities between that particular user and other items, and/or based on some number of known affinities between other users and that particular item or other items. The term affinity, as used throughout this disclosure, refers to a liking, attraction, preference, interest, etc., that a particular user may harbor for a particular item. The affinity may be determined based on information known about the user or people similar to the user or may be received from the user. The affinity may be a potential affinity that is predicted by the recommendation system 4400 or another entity. Instead, the affinity may be an actual or known affinity that is confirmed by the user or by user behavior (such as, for example, an affinity associated with a user and an item that is determined based on input received from the user in connection with the item). Once the recommendation system 4400 predicts an affinity between a particular user and a particular item, the system may make recommendation regarding the particular item to the particular user based on the predicted affinity.

[0249] According to various exemplary embodiments, the matrix generation module 4402 may access or receive affinity ratings information, and generate a ratings matrix based on such affinity ratings information. The ratings matrix may include various matrix values that correspond to affinity ratings, where each of the affinity ratings describe a degree of affinity between (or associated with) one of various users and one of various items. For example, each affinity rating in the ratings matrix may be a number from 0 to 1, where 0 represents a complete dislike of the item by the user (i.e., a low affinity between user and item), and a 1 represents that the user really likes the item (i.e., a high affinity between user and item).

[0250] The matrix generation module 4402 may access or receive affinity ratings information from, for example, the ratings database 4406. The affinity ratings information may identify various users and known affinities (e.g., likes or dislikes) each user may have for one or more items. Fig. 45 illustrates non-limiting examples of affinity ratings information the matrix generation module 4402 may receive or access. For example, table 4500 is an exemplary data structure that may store affinity ratings information indicating various product items (such as smartphone, tablet, lemonade, etc.), and a list of users that have a high affinity for that item (i.e., users that like that item), as well as a list of users that have a low affinity for that item (i.e., users that do not like that item). For example, as illustrated in table 4500, the smartphone item is liked by the user George, but is not liked by the user Miguel, and so forth. Table 4502 is an exemplary data structure that stores a different type of affinity ratings information indicating various users (such as George, Miguel, etc.), and a list of items for which that user has a high affinity (i.e., items that the user likes), as well as a list of items for which that user has a low affinity (i.e., items that the user does not like). For example, as illustrated in table 4502, the user George likes the smartphone item and tablet item, but does not like the scooter item, etc. As described in this disclosure, the term 'item' refers to any product, item, object, location, service, retail category, area of interest, activity, hobby, a webpage, an advertisement, media content, etc., that may be associated with a user.

[0251] Such affinity ratings information (and/or such tables or data structures storing such affinity ratings information) illustrated in Fig. 45 may be stored in, for example, the ratings database 4406 illustrated in Fig. 44. Such affinity ratings information may be obtained in various ways, and from various sources. For example, the affinity ratings information may reflect a high affinity rating associated with a user and an item, because that user purchased that item, searched for that item, submitted a request for that item, attempted to access information associated with that item, and so forth. As another example, when a system such as machine learning facility 1318 illustrated in Fig. 13 provides a decision to a user in 1310, the decision may be a recommendation for an item (such as an smartphone or a scooter), and the user may provide feedback to the machine learning facility 1318 in 1312. If the feedback indicates that the user

found the recommendation helpful, then the affinity ratings information associated with the user may reflect that the user liked the recommended item or has a high affinity for the recommended item. On the other hand, if the feedback indicates that the user did not find the recommendation helpful, then the affinity ratings information associated with the user may reflect that the user did not like the recommended item or that the user has a low affinity for the recommended item. Thus, the affinity ratings information may be generated by a machine learning facility (e.g., machine learning facility 1318 illustrated in Fig. 13) or recommendation system 4400.

[0252] While the exemplary affinity ratings information of Fig. 45 indicates that the user either liked or disliked the corresponding item, it should be understood that the affinity ratings information may reflect a range of affinity ratings associated with the particular user and the particular item. For example, the affinity ratings information may indicate a number from 0 to 1, where 0 represents a complete dislike of the item by the person, and a 1 represents that the person absolutely loves the item. Other schemes for indicating a degree of affinity harbored by a user may be utilized will be apparent to those skilled in the art.

[0253] After the matrix generation module 4402 receives or accesses the affinity ratings information, the matrix generation module 4402 generates a ratings matrix based on such affinity ratings information. For example, Fig. 46 illustrates an example of ratings matrix R 4600 generated based on the exemplary affinity ratings information of Fig. 45. Ratings matrix R is a $k \times j$ matrix (i.e., a matrix having k rows and j columns) and includes affinity ratings that each indicate a degree affinity associated with one of various users and one of various items. In particular, each of the k rows of ratings matrix R (i.e., rows $k1-k3$) identifies a particular user (e.g., George, Robert, etc.), and includes known affinity ratings associated that user (e.g., ratings $r11$, $r12$ and $r14$ associated with the user George). Each of the j columns of ratings matrix R (i.e., rows $j1-j4$) identifies a particular item (e.g., smartphone, tablet, etc.) and includes known affinity ratings associated that particular item (e.g., ratings $r23$ and $r33$ associated with the lemonade item). The affinity ratings (represented by $r11$, $r12$, etc.) may correspond to, for example, a number from 0 to 1, where 0 represents a complete dislike of the item by the user, and a 1 represents that the

user absolutely “loves” the item. Thus, any particular entry at a specific row k and column j in ratings matrix R is configured to include an affinity rating describing a degree of affinity associated with the user of the corresponding row k and the item of the corresponding column j , and wherein such an affinity rating indicates a degree of affinity that the particular user has for the particular item. For example, the affinity rating r_{23} at row k_2 and column j_3 in ratings matrix R is associated with the user Robert (of the corresponding row k_2) and associated with the item lemonade (of the corresponding column j_3), and wherein such an affinity rating r_{23} indicates a degree of affinity that Robert has for lemonade.

[0254] As illustrated in Fig. 46, many cells in the ratings matrix J may be missing, such as in the case when an affinity rating for a particular combination of person and item is unknown. For example, the affinity rating associated with user Miguel and the tablet item is not known, and thus the corresponding cell of the ratings matrix J (at row k_3 and column j_2) is missing. Thus, the ratings matrix R may include at least one missing entry representing an unknown affinity rating associated with a specific user and a specific item.

[0255] The affinity predication module 4404 is configured to find a new matrix R' (Fig. 49) that has no missing entries. That is, the prediction module 4404 is configured to receive ratings matrix R , which has one or more known affinity ratings as well as one or more missing affinity ratings, and generate a revised matrix R' that includes predicted affinity ratings in place of the missing affinity ratings. Thus, revised matrix R' includes either a known affinity rating or a predicted affinity rating for every combination of user and item.

[0256] To accomplish this, the affinity predication module 4404 is configured to produce a factorization of ratings matrix R into two complete factor matrices, referred to herein as a user matrix U and an item matrix I . For all known affinity ratings at index (x,y) in ratings matrix R , $R[x,y] \sim \text{dot}(U[x],I[y])$. Put another way, for every known affinity rating r at a particular row k and column j in the ratings matrix R , the dot product of the corresponding row k in user matrix U and corresponding column j in item matrix I produces this affinity rating r .

[0257] Fig. 47 illustrates an exemplary user matrix U 4700 that is a $k \times p$ matrix (i.e., a matrix having k rows and p columns) and includes various “user parameters” u_{11} - u_{35} , wherein the user parameters may be numbers, for

example. Each of the k rows of user matrix U (i.e., rows $k1$ - $k3$) is associated with and/or identifies a particular user (e.g., George, Robert, etc.), and includes various user parameters u associated with and/or describing that particular user. For example, row $k2$ identifies user Robert and includes the user parameters $u21$ - $u25$. Fig. 48 illustrates an exemplary item matrix I 4800 that is a $p \times j$ matrix (i.e., a matrix having p rows and j columns) and includes various "item parameters" $i11$ - $i54$, wherein the item parameters may be numbers, for example. Each of the j columns of item matrix I (i.e., columns $j1$ - $j4$) is associated with and/or identifies a particular item (e.g., smartphone, tablet, etc.), and includes various item parameters i associated with and/or describing that particular item. For example, row $j2$ identifies the tablet item and includes the item parameters $i12$, $i22$, $i32$, $i42$, and $i52$. Thus, for known affinity rating $r22$ at row $k2$ and column $j2$ in the ratings matrix R (see Fig. 46), the dot product of the corresponding row $k2$ in user matrix U (see Fig. 47) and corresponding column $j2$ in item matrix I (see Fig. 48) produces this affinity rating $r22$. That is, $(u21 \times i12) + (u22 \times i22) + (u23 \times i32) + (u24 \times i42) + (u25 \times i52) = r22$. Aspects of matrix multiplication are well understood by those skilled in the art and shall not be described further in the interests of clarity.

[0258] Moreover, if user matrix U and item matrix I are complete matrices with no missing values, then the revised ratings matrix $R' = UI$ (i.e., the matrix multiplication product of user matrix U and item matrix I), since $R' = UI$ will have a value for every combination of a user and item in ratings matrix R . Since matrix R has many missing values, there are a variety of ways for prediction module 4404 to factor matrix R into complete matrices U and I . One example apparent to those skilled in the art is the Singular Value Decomposition (SVD) technique. This technique, along with others, attempts to produce a factorization that minimizes some error (loss) function between R and R' , and there are many choices for the kernel function that decides the loss function.

[0259] In the SVD technique well known by those skilled in the art, the singular value decomposition of an $m \times n$ real or complex matrix M is a factorization of the form $M = U\Sigma V^*$, where U is an $m \times m$ real or complex unitary matrix, Σ is an $m \times n$ rectangular diagonal matrix with nonnegative real numbers on the diagonal, and V^* (the conjugate transpose of V) is an $n \times n$ real or complex unitary matrix. The prediction module 4404 is configured to factorize

ratings matrix R into user matrix U and item matrix I , and thereby generate revised ratings matrix $R' = UI$, based on the SVD technique.

[0260] Moreover, there are various ways to compute the SVD factorization, as understood by those skilled in the art. One example is computing the SVD by using stochastic gradient descent, while another is computing the SVD using Alternating Least Squares (ALS). Thus, prediction module 4404 may compute the SVD of ratings matrix R via stochastic gradient descent, computing one column of item matrix I at a time, where successive columns are computed. The factorization of a matrix using SVD via stochastic gradient descent is understood by those skilled in the art, and will not be described in further detail herein in the interests of brevity.

[0261] In the alternating least squares (ALS) technique, an initial choice of values for matrix U and matrix I are made, and then a series of updates is made, first to matrix U and then to matrix I , in order to produce an approximation to matrix R based on matrix R 's known values. The updates to matrix U are made by generating a least-squares fit solution of $UI = R$, solving for matrix U . Then the updates to matrix I are made by generation a least-squares fit solution of $UI = R$ solving for matrix I . This process is repeated until the magnitude of changes to matrix U and matrix I are below some threshold, until a fixed maximum number of steps has been reached, or until a variable number of steps is reached based on properties of matrix R . For example, if matrix R represents a graph the number of updates may be proportional to the diameter of the graph.

[0262] Note that the factors U and V^* produced by SVD are essentially the same as matrices U and I , except that the Σ matrix produced by SVD is multiplied into both U and I in some arbitrary way. As described above, SVD is essentially a theorem stating that a matrix can be uniquely factored into U , Σ , and V^* , and there are many implementations of algorithms for performing that factorization. A true SVD is only specified for dense matrices, and algorithms such as the lanczos algorithm are typically used to perform the operation. Given that the original ratings matrix R has missing entries, the recommendation system 4400 may probabilistically estimate what U and V^* (corresponding to user matrix U and item matrix I) might be, using algorithms like the well known alternating least squares algorithm/technique.

[0263] Thus, the prediction module 4404 is configured to factorize ratings matrix R into user matrix U and item matrix I , and thereby generate revised ratings matrix $R' = UI$, based on the ALS technique. Fig. 49 illustrates a revised ratings matrix R' 4900 generated by the prediction module 4404. Whereas the ratings matrix R illustrated in Fig. 45 had one or more missing affinity ratings, the revised ratings matrix R' includes predicted affinity ratings r_{13} , r_{21} , r_{24} and r_{32} in place of the missing affinity ratings. Thus, revised matrix R' includes either a known affinity rating or a predicted affinity rating for every combination of user and item. The factorization of a matrix based on ALS is understood by those skilled in the art, and will not be described in further detail herein in the interests of brevity.

[0264] Based on the predicted affinity ratings, the prediction module 4404 may provide a recommendation regarding a specific item to a specific user. For example, if a predicted affinity rating for a particular user and particular item is 1, then the module 4404 may recommend that item to the user, whereas if the predicted affinity rating is 0, then the module 4404 may not recommend that item to the user.

[0265] Fig. 50 is a flowchart illustrating an example method 5000, according to various embodiments. The method 5000 may be performed at least in part by, for example, the recommendation system 4400 illustrated in Fig. 44 (or an apparatus having similar modules). In 5001, the recommendation system 4400 generates a ratings matrix. An example of such a ratings matrix R is illustrated in Fig. 46. In 5002, the recommendation system 4400 factorizes ratings matrix R into user matrix U and item matrix I , and thereby generates revised ratings matrix $R' = UI$, using the aforementioned ALS technique, for example. The revised ratings matrix R' includes at least one entry representing a predicted affinity rating describing a degree of affinity associated with a specific user and a specific item, and an example of such a revised ratings matrix R' is illustrated in Fig. 49. In 5003, the recommendation system 4400 provides a recommendation regarding a specific item to a specific user, based on the predicted affinity rating in revised ratings matrix R' associated with the specific user and the specific item.

[0266] The ALS process performed by the prediction module 4404 - for factoring the ratings matrix R into user matrix U and item matrix I to generate

revised ratings matrix R' - may take a significant amount of time on current computer systems. Accordingly, there is a potential problem when a new user or new item is to be added to the ratings matrix R , in that it may be infeasible to quickly add the new user or the new item to the ratings matrix R , and then proceed through the entire ALS process again to factor the new ratings matrix R and to generate a new revised ratings matrix R' .

[0267] Accordingly, in various exemplary embodiments, the prediction module 4404 may perform an "incremental approximate update" to add a single new row or column to ratings matrix R without re-computing all of user matrix U and item matrix I . For example, when a new user needs to be added to the system (as happens when a new users signs up at a website or uses a mobile recommendation application, such as those described in various exemplary embodiments of this disclosure), the prediction module 4404 may estimate the user parameters u of the new user for direct insertion into the user matrix U . In particular, the prediction module 4404 may estimate the user parameters u of the new user by generating a solution (e.g., ALS solution) of $Iu = r$ by solving for u , where r are the new user's ratings of some of the existing j items in ratings matrix R .

[0268] For example, Fig. 51 illustrates an example of the same ratings matrix R illustrated in Fig. 46, except that a new row $k4$ corresponding to a new user Sven has been inserted into the ratings matrix R by the prediction module 4404. The row $k4$ identifying the new user Sven includes various known affinity ratings (e.g., $r41$, $r42$ and $r44$), where these known ratings describe a degree of affinity associated with the new user Sven and associated with various existing items in columns $j1$ - $j4$. For example, rating $r41$ describes a degree of affinity associated with the new user Sven and the smartphone item, rating $r44$ describes a degree of affinity associated with the new user Sven and the scooter item, and so forth. Fig. 52 illustrates an example of user matrix U with user parameters $u41$ - $u45$ describing the new user Sven that may be estimated by the prediction module 4404.

[0269] In particular, the prediction module 4404 generates a solution (e.g., ALS solution) of $Iu = r$ by solving for u . That is, since the item matrix I with item parameters i describing each of the items has already been generated (see Fig. 48), the prediction module 4404 may assume that the product of user

parameters u_{41} - u_{45} describing the new user Sven and the item parameters corresponding to a specific item should produce the known affinity rating describing the degree of affinity associated with the user Sven and that specific item. For example, prediction module 4404 may assume that the product of user parameters u_{41} - u_{45} describing the new user Sven (see Fig. 52) and the item parameters $[i_{14}, i_{24}, i_{34}, i_{44}, i_{54}]$ describing the scooter item (see Fig. 48) will produce the known affinity rating r_{44} describing the degree of affinity associated with the user Sven and the scooter item (see Fig. 51). Thus, based on the already generated item matrix I (and item parameters i therein) and based on all the known affinity ratings r associated with the new user, the prediction module 4404 generates a solution (e.g., ALS solution) of $Iu = r$ by solving for u , in order to estimate the user parameters u of the new user (e.g., u_{41} - u_{45} illustrated in Fig. 52) and insert them into the user matrix U (see Fig. 52).

[0270] Moreover, according to various exemplary embodiments, once the user parameters of the new user are estimated, these user parameters may be utilized by the prediction module 4404 to estimate missing affinity ratings associated with the new user. For example, the prediction module 4404 may generate the dot product of the estimated user parameters $[u_{41}$ - $u_{45}]$ describing user Sven (see Fig. 52) and the item parameters $[i_{13}, i_{23}, i_{33}, i_{43}$ and $i_{53}]$ describing the lemonade item (see Fig. 48), to produce an estimate of the missing affinity rating describing a degree of affinity associated with the user Sven and the lemonade item (see Fig. 51).

[0271] Similarly, when a new item needs to be added to the system, the prediction module 4404 in various embodiments may estimate the item parameters i of the new item for direct insertion into the item matrix I , by generating a solution (e.g., ALS solution) of $U_i = u$ by solving for i , where u are ratings of the new item by some of the existing k users in ratings matrix R .

[0272] For example, Fig. 53 illustrates an example of the same ratings matrix R illustrated in Fig. 46, except that a new column j_5 corresponding to a new printer item has been inserted into the ratings matrix R by the prediction module 4404. The row j_5 identifying the new printer item includes various known affinity ratings (e.g., r_{15}, r_{35}) describing a degree of affinity associated with the new printer item and associated with various existing users in columns k_1 - k_3 . For example, rating r_{15} describes a degree of affinity associated with the

new printer item and the user George, rating r_{35} describes a degree of affinity associated with the new printer item and the user Miguel, etc. Fig. 54 illustrates an example of item matrix I with item parameters i_{15} , i_{25} , i_{35} , i_{45} and i_{55} describing the new printer item that may be estimated by the prediction module 4404.

[0273] In particular, the prediction module 4404 generates a solution (e.g., ALS solution) of $U_i = r$ by solving for i , where r are ratings of the new item by some of the existing k users in ratings matrix R . That is, since the user matrix U with user parameters u describing each of the users has already been generated (see Fig. 47), the prediction module 4404 may assume that the product of item parameters i describing the new item and the user parameters corresponding to a specific user will produce the known affinity rating describing the degree of affinity associated with the new item and that specific user. For example, prediction module 4404 may assume that the product of $[i_{15}, i_{25}, i_{35}, i_{45} \text{ and } i_{55}]$ describing the new printer item (see Fig. 54) and the user parameters $[u_{11}-u_{15}]$ describing the user George (see Fig. 47) will produce the known affinity rating r_{15} describing the degree of affinity associated with the printer item and the user George (see Fig. 53). Thus, based on the already generated user matrix U (and user parameters u therein) and based on all the known affinity ratings r associated with the new item, the prediction module 4404 generates a solution (e.g., ALS solution) of $U_i = r$ by solving for i , in order to estimate the item parameters i of the new item (e.g., i_{15} , i_{25} , i_{35} , i_{45} and i_{55} illustrated in Fig. 54) and insert them into the item matrix I (see Fig. 54).

[0274] Moreover, according to various exemplary embodiments, once the item parameters of the new item are estimated, these item parameters may be utilized by the prediction module 4404 to estimate missing ratings associated with the new item. For example, the prediction module 4404 may generate the dot product of the estimated item parameters $[i_{15}, i_{25}, i_{35}, i_{45} \text{ and } i_{55}]$ for the printer item (see Fig. 54) and the user parameters $[u_{21}-u_{25}]$ describing the user Robert (see Fig. 47), to produce an estimate of the missing affinity rating describing a degree of affinity associated with the printer item and the user Robert (see Fig. 53).

[0275] Fig. 55 is a flowchart illustrating an example method 5500, according to various embodiments. The method 5500 may be performed at least

in part by, for example, the recommendation system 4400 illustrated in Fig. 44 (or an apparatus having similar modules). In 5501, the system inserts a new row identifying a new user into the ratings matrix, the new row including a known affinity rating describing a degree of affinity associated with the new user. In 5502, the system estimates user parameters describing the new user, based on the known affinity rating describing the degree of affinity associated with the new user and item parameters included in the item matrix. In 5503, the system predicts an affinity rating describing a degree of affinity associated with the new user, based on the estimated user parameters and item parameters included in the item matrix.

[0276] Fig. 56 is a flowchart illustrating an example method 5600, according to various embodiments. The method 5600 may be performed at least in part by, for example, the recommendation system 4400 illustrated in Fig. 44 (or an apparatus having similar modules). In 5601, the system inserts a new column describing a new item into the ratings matrix, the new column including a known affinity rating describing a degree of affinity associated with the new item. In 5601, the system estimates item parameters describing the new item, based on the known affinity rating describing the degree of affinity associated with the new item and user parameters included in the user matrix. In 5601, the system predicts an affinity rating describing a degree of affinity associated with the new item, based on the estimated item parameters and user parameters included in the user matrix. A further natural extension allows r or u to contain a mix of ratings of/for existing users/items in R with new users and items not in R .

[0277] According to various exemplary embodiments, a similar technique to the aforementioned “incremental approximate update” may also be performed by the prediction module 4404 to update user parameters u and item parameters i , based on new affinity ratings information (associated with existing users and items) that is received after ratings matrix R is already factored into user matrix U and item matrix I .

[0278] For example, when an existing user represented by a row k in ratings matrix R provides a new rating of an existing product represented by a column j , that new rating can be used to update the user parameters u of that user immediately, rather than waiting until a new complete factorization of ratings matrix R is complete. For example, referring to the ratings matrix R in Fig. 46,

suppose a new rating r_{31} describing a degree of affinity associated with existing user Miguel (represented by a row k_3) and the smartphone item (represented by column j_1) is received.

[0279] The prediction module 4404 generates a solution (e.g., ALS solution) of $Iu = r$ by solving for u . That is, since the item matrix I with item parameters i describing each of the items has already been generated (see Fig. 48), the prediction module 4404 may assume that the product of user parameters describing an existing user and the item parameters describing an existing item should produce the known affinity rating describing the degree of affinity that is associated with the existing user and that existing item. For example, prediction module 4404 may assume that the product of user parameters $[u_{31}-u_{35}]$ describing the user Miguel (see Fig. 47) and the item parameters $[i_{11}, i_{21}, i_{31}, i_{41}, i_{51}]$ describing the smartphone item (see Fig. 48) should produce the new affinity rating r_{31} describing the degree of affinity associated with the user Miguel and the smartphone item. Thus, based on the already generated item matrix I (and item parameters i therein) and based on all the known affinity ratings r associated with a user (including the newly received affinity rating), the prediction module 4404 generates a solution (e.g., ALS solution) of $Iu = r$ by solving for u , in order to revise the user parameters u in the user matrix U .

[0280] Instead or in addition, when an existing user represented by a row k in ratings matrix R provides a new rating of an existing product represented by a column j , that new rating can be used to update the item parameters i of that existing product immediately, rather than waiting until a new complete factorization of ratings matrix R is complete. For example, referring to the ratings matrix R in Fig. 46, suppose a new rating r_{31} describing a degree of affinity associated with existing user Miguel (represented by a row k_3) and the smartphone item (represented by column j_1) is received.

[0281] The prediction module 4404 generates a solution (e.g., ALS solution) of $Ui = r$ by solving for i . That is, since the user matrix U with user parameters u describing each of the users has already been generated (see Fig. 47), the prediction module 4404 may assume that the product of item parameters describing an existing item and the user parameters describing an existing user should produce the known affinity rating describing the degree of affinity associated with that existing item and user. For example, prediction module

4404 may assume that the product of item parameters $[i11, i21, i31, i41, i51]$ describing the existing smartphone item (see Fig. 48) and the user parameters $[u31-u35]$ describing the user Miguel (see Fig. 47) should produce the new affinity rating $r31$ describing the degree of affinity associated with the smartphone and the user Miguel. Thus, based on the already generated user matrix U (and user parameters u therein) and based on all the known affinity ratings r associated with an item (including the newly received affinity rating), the prediction module 4404 generates a solution (e.g., ALS solution) of $U_i = r$ by solving for i , in order to revise the item parameters i in the item matrix I .

[0282] Moreover, according to various exemplary embodiments, after the prediction module 4404 updates the parameters for a user submitting a new rating, the prediction module 4404 updates the item parameters for other items rated by that user. For example, since the dot product of the user parameters for a particular user and the item parameters for a particular item should produce the known affinity ratings describing the degree of affinity associated with the user and the particular item, the updating of the user parameters for a particular user (e.g. user parameters $u31-u35$ for user Miguel in Fig. 47) may require updates to the item parameters for various items (e.g. the lemonade item and the scooter item in Fig. 48), given the known affinity rating $r33$ describing the degree of affinity associated with the user Miguel and the lemonade item, and given known affinity rating $r34$ describing the degree of affinity associated with the user Miguel and the scooter item (see Fig. 46). That is, the dot product of the revised user parameters $[u31-u35]$ describing Miguel and the item parameters $[i14, 24, 34, 44, 54]$ describing the scooter item should equal the affinity rating $r34$ describing the degree of affinity associated with the user Miguel and the scooter item (see Fig. 46). Thus, based on the revised user parameters u describing a user and all the known affinity ratings r associated with the user (including the newly received affinity rating), the prediction module 4404 generates a solution (e.g., ALS solution) of $U_i = r$ by solving for i , in order to revise the item parameters i in the item matrix I .

[0283] Similarly, according to various exemplary embodiments, after the prediction module 4404 updates the item parameters for an item having a new rating, the prediction module 4404 updates the user parameters for other users that rated that item. For example, since the dot product of the item parameters

for a particular item and the user parameters for a particular user should produce the known affinity ratings describing the degree of affinity associated with the particular user and the particular item, the updating of the item parameters for a particular item (e.g., $[i11, 21, 31, 41, 51]$ describing the smartphone item in Fig. 48) may require updates to the user parameters for various users (e.g. the user George in Fig. 47), given the known affinity rating $r11$ describing the degree of affinity associated with the smartphone item and the user George. That is, the dot product of the revised item parameters $[i11, 21, 31, 41, 51]$ describing the smartphone item and the user parameters $[u11-u15]$ describing the user George must equal the affinity rating $r11$ describing the degree of affinity associated with the user George and the smartphone item (see Fig. 46). Thus, based on the revised item parameters i describing an item and all the known affinity ratings r associated with the item (including the newly received affinity rating), the prediction module 4404 generates a least squares solution (e.g., ALS solution) of $Iu = r$ by solving for u , in order to revise the user parameters u in the user matrix U .

[0284] Fig. 57 is a flowchart illustrating an example method 5700, according to various embodiments. The method 5700 may be performed at least in part by, for example, the recommendation system 4400 illustrated in Fig. 44 (or an apparatus having similar modules). In 5701, the system receives a new affinity rating describing a degree of affinity associated with a particular user and a particular item. In 5702, the system revises user parameters in the user matrix that describe the particular user, based on the new affinity rating. In 5703, the system revises item parameters for other items rated by that particular user, based on the revised user parameters of the particular user.

[0285] Fig. 58 is a flowchart illustrating an example method 5800, according to various embodiments. The method 5800 may be performed at least in part by, for example, the recommendation system 4400 illustrated in Fig. 44 (or an apparatus having similar modules). In 5801, the system receives a new affinity rating describing a degree of affinity associated with a particular user and a particular item. In 5802, the system revises item parameters in the item matrix that describe the particular item, based on the new affinity rating. In 5803, the system revises user parameters for other users that have rated that particular item, based on the revised item parameters of the particular item.

[0286] The aforementioned incremental process for computing parameters for new users and items, or updating parameters for existing users and items, may be used to first compute the factorization of R instead of relying on an offline/batch factorization of R ahead of time. The incremental process of the aforementioned embodiments can be used with a factorization of R created by SVD based on ALS techniques.

[0287] A common problem with computing parameters for new users or items may be that it requires an existing set of ratings for the new user or users that have rated the item for new items. One technique for dealing with this is to "cold start" new items by training them with a set of tags that are treated as if they were users. These artificial "users" in turn are trained by existing items that have the same tags. A similar approach can be used to cold start new users – training them with fake items each of which is represented by a tag that other existing users have tagged themselves with.

[0288] According to various exemplary embodiments, the system 440 incorporates the concept of confidence into the ratings of the ratings matrix R , by generating a confidence matrix C with the same dimensions of ratings matrix R , where every non-missing entry (i.e. every known affinity rating) in ratings matrix R has a corresponding entry in confidence matrix C that specifies the confidence of the corresponding value in ratings matrix R .

[0289] The prediction module 4404 may estimate a confidence value associated with every known affinity rating in the ratings matrix R based on various factors. As some non-limiting examples, the prediction module 4404 may take into account (1) whether a positive affinity rating associated with a user and item was generated based on the user clicking on a page associated with the item (resulting in a low confidence value), or (2) whether a positive affinity rating associated with a user and item was generated based on the user providing positive feedback regarding the item to a machine learning facility, such as machine learning facility 1318 in Fig. 13 (resulting in a high confidence value), or (3) whether a positive affinity rating associated with a user and item was generated based on the user purchasing the item (resulting in a high confidence value), or (4) based whether the user has not returned the purchased item after a predetermined time period (resulting in a higher confidence value), and so on.

The confidence value may, for example, be a number between 0 and 1, 0 indicating no confidence and 1 indicating the highest confidence.

[0290] After the prediction module 4404 estimates a confidence value associated with every known affinity rating in the ratings matrix R , the prediction module 4404 may generate the confidence matrix C that is substantially similar to the ratings matrix R , except the confidence matrix C includes the confidence values corresponding to each rating associated with a particular user and item. Fig. 59 illustrates an example of a confidence matrix C 5900 that is similar to the ratings matrix R illustrated in Fig. 46, except the confidence matrix C includes confidence values corresponding to the various ratings. For example, the confidence value c_{31} in confidence matrix C indicates the confidence associated with the affinity rating r_{31} (see Fig. 46) describing a degree of affinity associated with the user Miguel and the smartphone item.

[0291] Thus, in an online shopping application the prediction module 4404 may store an affinity rating in $R(k,j) = 1.0$ when a user identified in column k clicks on a product listing page for an item identified in column j , indicating a high affinity between the user and the item. Nevertheless, the prediction module 4404 may store a very low confidence $C(k,j) = 0.1$, indicating that while the user may like the item since they clicked on the item's product page, it is a very weak indicator of affinity. As another example, if the same user buys the same item, the prediction module 4404 may store an affinity rating of $R(k,j) = 1.0$, indicating a high affinity between the user and the item, and the prediction module 4404 may store a very high confidence $C(k,j) = 0.9$ indicating that it is a strong indicator of affinity. Further, if the user has not returned to the store after two weeks, prediction module 4404 may increase the confidence value to $C(k,j) = 0.99$.

[0292] Given the confidence values associated with each of the affinity ratings in matrix R , the prediction module 4404 may utilize a weighted least squares approach when factorizing the ratings matrix R into user matrix U and item matrix I in order to generate revised matrix R' . That is, instead of the prediction module 4404 computing the user matrix U and item matrix I using an alternating least-squares approach where all of the affinity ratings in ratings matrix R are given the same weight, each of the affinity ratings in ratings matrix R may be weighted by the prediction module 4404 based on the corresponding

confidence value in confidence matrix C , such that the ALS technique used by the prediction module 4404 produces a better recommendation result. Factorization of matrices using a weighted least squares approach is well understood by those skilled in the art, and shall not be described in more detail herein in the interests of brevity.

[0293] According to various exemplary embodiments, the recommendation system 4400 may further improve in recommendation accuracy by estimating whether missing affinity ratings in ratings matrix R result from a user not being interested in the items they didn't rate or whether the user was unfamiliar with the items they didn't rate. For example, a man may not rate many women's shoes he views on a shopping web site because he's not interested in women's shoes. On the other hand, he may not have rated a new restaurant simply because he has never heard of it, even if after going to it he would like it.

[0294] According to one approach, the prediction module 4400 may address this by randomly sampling items that have not been rated by each user and assume a low confidence dislike between the user and the item. This assumes that there is a reason that the user has not rated the item, such as not being interested in the item.

[0295] According to another approach, the prediction module 4400 may address this by first estimating the probability $p(k,j)$ that a user k has had a chance or opportunity to rate item j (or follow, purchase, click, etc. the item j), and then assigning a missing affinity rating $R(k,j)$ in matrix R a value close to 0 (meaning dislike or low affinity between the user k and item j) with confidence value of $C(k,j) = p(k,j)$. This indicates that an item a user has had a chance to rate but has not rated, should be considered with high confidence to be an item the user is not interested in, and so shouldn't be recommended to that user. This also indicates that an item a user has not had a chance to rate should be considered with little confidence to be an item the user is not interested in. In the example above, the user had ample opportunity to search out and rate women's shoes, but never did, and so the prediction module 4400 assumes that the user is not interested in women's shoes. On the other hand, the user did not have a chance to see the new restaurant, and so the prediction module 4400 does not assume the new restaurant is disliked by the user.

[0296] Fig. 60 is a flowchart illustrating an example method 6000, according to various embodiments. The method 6000 may be performed at least in part by, for example, the recommendation system 4400 illustrated in Fig. 44 (or an apparatus having similar modules). In 6001, the recommendation system 4400 generates a ratings matrix. An example of such a ratings matrix R is illustrated in Fig. 46. In 6002, the system may estimate a confidence value associated with every known affinity rating in the ratings matrix R. In 6003, the system may generate a confidence matrix C that is substantially similar to the ratings matrix R, except the confidence matrix C includes the confidence values corresponding to each rating associated with a particular user and item. Fig. 59 illustrates an example of a confidence matrix C 5900 similar to the ratings matrix R illustrated in Fig. 46, except the confidence matrix C includes confidence values corresponding to the various affinity ratings.

[0297] In 6004, the system estimates the probability that a user has had a chance to rate an item, where the user and the item are associated with a missing affinity rating in ratings matrix R. In 6005, the system replaces the missing affinity rating in ratings matrix R with a low affinity rating (e.g. 0). In 6006, the system inserts a confidence value associated with the user and the item in a confidence matrix C, where the confidence value has a value equal to the probability estimated in 6004. In 6007, the system generates a revised ratings matrix R', by factoring the ratings matrix R into user matrix U and item matrix I, based on a weighted least squares technique that takes into account the confidence values (i.e., that weights each affinity rating in the ratings matrix based on the corresponding confidence value in the confidence matrix).

[0298] According to various exemplary embodiments, the recommendation system 4400 is configured to deal interpret non-negative feedback. For example, in systems like Facebook's ® "Like" button or Google's ® "+ 1" button there is no corresponding dislike button. This means that there can be no negative feedback, only positive feedback on which items a user likes. Similar unbalanced rating problems arise in: social networks (e.g., where you can only follow people); online/offline shopping (where the vast majority of data, such as purchases, clicks, products being added to shopping carts, etc., only shows positive feedback/ratings); search engine result selection; advertising optimization, and so forth. Accordingly, it may be difficult to interpret whether

missing affinity values in the ratings matrix R correspond to negative feedback. The aforementioned embodiments of modeling a probability of user u having had the opportunity to rate, follow, purchase, click, and the like on item, can be used to model missing values in R as negative feedback, which may provide significant increase in recommendation accuracy.

[0299] For example, according to various exemplary embodiments, the recommendation system 4400 estimates a probability in $p(j,k)$ that a missing affinity rating corresponds to a negative affinity rating. For example, the prediction module 4404 may organize all the k items in ratings matrix R into a taxonomy of different categories t , such as the smartphone item and tablet item being in the category 'electronics', the lemonade item being in the category 'food and beverages', the scooter item being in the category 'vehicles and transportation', and so forth. Thereafter, the prediction module 4404 may assume $p(k,j) \sim p(k,t)$ where item j is in category t . Note that $p(k,j)$ is the probability that a missing affinity rating associated with user k and item j corresponds to a negative affinity rating, and $p(k,t)$ is the probability that a missing affinity rating associated with user k and an item in category t corresponds to a negative affinity rating.

[0300] The value of $p(k,t)$ itself may be computed prediction module 4404 as the fraction of user j 's ratings that are associated with category t . For example, referring to Fig. 46, since half of user Robert's known ratings are in the category of electronics (one rating r_{22} is for the tablet device, the other rating r_{23} is for lemonade), then $p(k,t)$ - or the probability that Robert's missing affinity rating in the category of electronics corresponds to negative feedback - can be calculated by prediction module 4404 to be 0.5. As another example, if all of a user k 's known ratings are in a single category t , then $p(k,t)$ - or the probability that a missing affinity rating in the category t corresponds to negative feedback - can be calculated by prediction module 4404 to be 1.

[0301] Alternatively, $p(k,t)$ could be estimated as what fraction of items in category t user k has rated. Alternatively, if each item j corresponds to a web page or advertisement, $p(k,j)$ may be computed based on what fraction of user k 's page views were of the page, advertisement or other content associated with item j . Alternatively, if item j is a search result, $p(k,j)$ can be computed as how

many times user k viewed item j , optionally scaled by the ranking of item j in the search results.

[0302] After the value of $p(k,j) \sim p(k,t)$ is determined, the prediction module 4400 may then assign a missing affinity rating $R(k,j)$ in matrix R a value close to 0 (meaning dislike or negative affinity between the user k and item j) with confidence value of $C(k,j) = p(k,j)$. This indicates that if the user has not rated a given item, but the user has frequently rated similar items, then the given item should be considered with high confidence to be an item the user is not interested in, and so shouldn't be recommended to that user. This also indicates that if the user has not rated a given item, and the user has rarely rated similar items, then the given item should be considered with low confidence to be an item the user is not interested in.

[0303] Fig. 61 is a flowchart illustrating an example method 6100, according to various embodiments. The method 6100 may be performed at least in part by, for example, the recommendation system 4400 illustrated in Fig. 44 (or an apparatus having similar modules). The method 6100 may be performed after 6003 in method 6000. In 6101, the system estimates the probability that a missing affinity rating in ratings matrix R that is associated with a user and an item corresponds to a negative affinity rating. In 6102, the system replaces the missing affinity rating in ratings matrix R with a low affinity rating (e.g. 0). In 6103, the system inserts a confidence value associated with the user and the item in a confidence matrix C , where the confidence value has a value equal to the probability estimated in 6101. In 6104, the system generates a revised ratings matrix R' , by factoring the ratings matrix R into user matrix U and item matrix I , based on a weighted least squares technique that takes into account the confidence values (i.e., that weights each affinity rating in the ratings matrix R based on the corresponding confidence value in the confidence matrix).

[0304] Another way for prediction module 4404 to estimate $p(u,i)$ may be to start create a matrix W which is a k by j matrix representing which users have been known to see or rate which items in the system. For example, a specific entry $W(k,j)$ may represent that the user k has not viewed or seen the item j , if the value of $W(k,j) = 0$, whereas the entry may represent that the user k has viewed or seen the item j , if the value of $W(k,j) = 1$. Reducing this matrix to a lower rank approximation may allow a better inference of which users have

been exposed or acted on which items. For example, computing the SVD or principle components of W allows the use of some number of those principle components to predict a new W' matrix that provides estimates for which items users have seen (but chose not to rate) than the matrix W does.

[0305] As described above, each row of the user matrix U includes user parameters associated with each user in the ratings matrix R , and each column of the item matrix I includes item parameters associated with each of the items in the ratings matrix R .

[0306] According to various exemplary embodiments, the recommendation system 4400 is configured to determine the user-to-user similarity of two or more users (e.g., determine whether a first user is similar to a second user), based on the user parameters in the user matrix U . For example, the prediction module 4404 may determine that a first user is similar to a second user, by determining the cosine similarity between two vectors corresponding to the user parameters of the first and second users, and determining that the cosine similarity is greater than or less than a specific threshold. For example, the prediction module 4404 may determine the similarity of the users Robert and Miguel by calculating the cosine similarity of the vectors $\{u_{21}, u_{22}, u_{23}, u_{24}, u_{25}\}$ (which corresponds to the user parameters for Robert, see Fig. 47) and $\{u_{31}, u_{32}, u_{33}, u_{34}, u_{35}\}$ (which corresponds to the user parameters for Miguel, see Fig. 47), and then determining whether the thus calculated cosine similarity is greater than or less than a predetermined threshold. The determination of the cosine similarity, which is a measure of similarity between two vectors by measuring the cosine of the angle between them, is well understood by those skilled in the art, and shall not be described in more detail herein in the interests of brevity.

[0307] As another example, the prediction module 4404 may determine that a first user is similar to a second user, by determining the Euclidean distance between two vectors corresponding to the user parameters of the first and second users, and determining that the calculated Euclidean distance is greater than or less than a specific threshold. The determination of Euclidean distance between vectors is well understood by those skilled in the art, and shall not be described in more detail herein in the interests of brevity. As yet another example, the prediction module 4404 may determine that a first user is similar to a second

user, by calculating the dot product of two vectors corresponding to the user parameters associated with the first and second users, and determining that the dot product is greater than or less than a specific threshold.

[0308] Similarly, the prediction module 4404 may determine the item-to-item similarity of two or more items (e.g., determine whether a first item is similar to a second item), based on the item parameters in the user matrix *I*, using Cosine similarity, Euclidean distance, dot product, etc., as described above with regards to determining user-to-user similarity. The aforementioned examples of cosine similarity, Euclidean distance and dot product are merely exemplary, and it should be understood that any other metrics and methods of measuring similarity between two vectors may be utilized by the prediction module 4404 to determine similarity between users and items.

[0309] After the prediction module 4404 determines that a first user is similar to a second user, the prediction module 4404 assigns recommendations or affinity ratings associated with the first user to the second user, and vice versa. For example, with reference to Fig. 46, if the users George and Miguel are similar, then George's affinity rating *r*₁₂ for the tablet may also be imputed to Miguel. If the system 4404 or a machine learning facility (such as machine learning facility 1318 illustrated in Fig. 13) decides to recommend an item to George, the system may recommend the item (or a similar item) to Miguel. Likewise, if the system 4404 or a machine learning facility decides not to recommend an item to George, the system may not recommend the item (or a similar item) to Miguel. Moreover, when the system 4404 or machine learning facility provides the recommendation to Miguel, the system or facility may inform the user Miguel that they are considered similar to the user George, and this item was also recommended to George (or George liked this item, etc.). Thus, user-to-user similarity may be important for showing who else is like a particular user and providing justification for why a recommendation is right. For example, the system 4404 or facility might not only show a particular user that a camera is a good recommendation, but may also show other people that like the camera who are similar to the particular user.

[0310] Similarly, after the prediction module 4404 determines that a first item is similar to a second item, the system assigns recommendations or affinity ratings associated with the first item to the second item, and vice versa. For

example, if the smartphone and tablet items are similar, then Miguel's affinity rating r_{31} for the smartphone (see Fig. 46) may also be imputed to the tablet. If the system 4404 or a machine learning facility (such as machine learning facility 1318 illustrated in Fig. 13) decides to recommend an item (e.g., smartphone) to a particular user (e.g., Miguel), the system may recommend a similar item (e.g., tablet) to that same user (e.g., Miguel).

[0311] Fig. 62 is a flowchart illustrating an example method 6200, according to various embodiments. The method 6200 may be performed at least in part by, for example, the recommendation system 4400 illustrated in Fig. 44 (or an apparatus having similar modules). In 6201, the system determines that a first user is similar to a second user, based on the user parameters in the user matrix that are associated with the first and second users. In 6202, the system generates a recommendation of a particular item for the second user (or generates an affinity rating describing a degree of affinity associated with the particular item and the first user), based on a recommendation of the particular item for the first user (or based on an affinity rating describing a degree of affinity associated with the particular item and the first user).

[0312] Fig. 63 is a flowchart illustrating an example method 6300, according to various embodiments. The method 6300 may be performed at least in part by, for example, the recommendation system 4400 illustrated in Fig. 44 (or an apparatus having similar modules). In 6301, the system determines that a first item is similar to a second item, based on the item parameters in the item matrix that are associated with the first and second items. In 6302, the system generates a recommendation of the second item for the particular user (or generates an affinity rating describing a degree of affinity associated with the second item and the particular user), based on a recommendation of the first item for the particular user (or based on an affinity rating describing a degree of affinity associated with the first item and the particular user).

[0313] According to various exemplary embodiments, the prediction module 4404 may determine user-to-user similarity by counting the number of affinity ratings a pair of users have in common (i.e., by counting the number of times a first user and a second user have a given affinity for the same item), and if the number of common ratings is greater than a predetermined threshold, the users may be considered to be similar by the prediction module 4404.

[0314] Moreover, the prediction module 4404 may take into account item-to-item similarity when counting the number of common ratings between two users (i.e., by counting the number of times a first user and a second user have a given affinity for similar items). For example, if Miguel has submitted a “like” rating of the smartphone and Robert has submitted a “like” rating of the tablet, and the smartphone and the tablet are considered to be similar items or have a high item-to-item similarity, then the prediction module 4404 may count this as a common rating of Miguel and Robert (i.e., both of them have something in common that they like). If the number of common ratings is greater than a predetermined threshold, the users may be considered to be similar by the prediction module 4404. After the prediction module 4404 determines that a first user is similar to a second user, the prediction module 4404 may assign recommendations or affinity ratings associated with the first user to the second user, and vice versa, as described in various embodiments above.

[0315] Fig. 64 is a flowchart illustrating an example method 6400, according to various embodiments. The method 6400 may be performed at least in part by, for example, the recommendation system 4400 illustrated in Fig. 44 (or an apparatus having similar modules). In 6401, the system counts a number of common affinity ratings between a first user and a second user, wherein each of the common affinity ratings are associated with a first item or a second item that is determined to be similar to the first item. In 6402, the system determines that the first user is similar to the second user, based on the number of common affinity ratings. In 6403, the system generates a recommendation associated with a particular item and the second user, based on a recommendation or affinity rating associated with the particular item and the first user.

MODULES, COMPONENTS AND LOGIC

[0316] Certain embodiments are described herein as including logic or a number of components, modules, or mechanisms. Modules may constitute either software modules (e.g., code embodied (1) on a non-transitory machine-readable medium or (2) in a transmission signal) or hardware-implemented modules. A hardware-implemented module is a tangible unit capable of performing certain operations and may be configured or arranged in a certain manner. In example embodiments, one or more computer systems (e.g., a

standalone, client or server computer system) or one or more processors may be configured by software (e.g., an application or application portion) as a hardware-implemented module that operates to perform certain operations as described herein.

[0317] In various embodiments, a hardware-implemented module may be implemented mechanically or electronically. For example, a hardware-implemented module may comprise dedicated circuitry or logic that is permanently configured (e.g., as a special-purpose processor, such as a field programmable gate array (FPGA) or an application-specific integrated circuit (ASIC)) to perform certain operations. A hardware-implemented module may also comprise programmable logic or circuitry (e.g., as encompassed within a general-purpose processor or other programmable processor) that is temporarily configured by software to perform certain operations. It will be appreciated that the decision to implement a hardware-implemented module mechanically, in dedicated and permanently configured circuitry, or in temporarily configured circuitry (e.g., configured by software) may be driven by cost and time considerations.

[0318] Accordingly, the term "hardware-implemented module" should be understood to encompass a tangible entity, be that an entity that is physically constructed, permanently configured (e.g., hardwired) or temporarily or transitorily configured (e.g., programmed) to operate in a certain manner and/or to perform certain operations described herein. Considering embodiments in which hardware-implemented modules are temporarily configured (e.g., programmed), each of the hardware-implemented modules need not be configured or instantiated at any one instance in time. For example, where the hardware-implemented modules comprise a general-purpose processor configured using software, the general-purpose processor may be configured as respective different hardware-implemented modules at different times. Software may accordingly configure a processor, for example, to constitute a particular hardware-implemented module at one instance of time and to constitute a different hardware-implemented module at a different instance of time.

[0319] Hardware-implemented modules can provide information to, and receive information from, other hardware-implemented modules. Accordingly, the described hardware-implemented modules may be regarded as being

communicatively coupled. Where multiple of such hardware-implemented modules exist contemporaneously, communications may be achieved through signal transmission (e.g., over appropriate circuits and buses) that connect the hardware-implemented modules. In embodiments in which multiple hardware-implemented modules are configured or instantiated at different times, communications between such hardware-implemented modules may be achieved, for example, through the storage and retrieval of information in memory structures to which the multiple hardware-implemented modules have access. For example, one hardware-implemented module may perform an operation, and store the output of that operation in a memory device to which it is communicatively coupled. A further hardware-implemented module may then, at a later time, access the memory device to retrieve and process the stored output. Hardware-implemented modules may also initiate communications with input or output devices, and can operate on a resource (e.g., a collection of information).

[0320] The various operations of example methods described herein may be performed, at least partially, by one or more processors that are temporarily configured (e.g., by software) or permanently configured to perform the relevant operations. Whether temporarily or permanently configured, such processors may constitute processor-implemented modules that operate to perform one or more operations or functions. The modules referred to herein may, in some example embodiments, comprise processor-implemented modules.

[0321] Similarly, the methods described herein may be at least partially processor-implemented. For example, at least some of the operations of a method may be performed by one or processors or processor-implemented modules. The performance of certain of the operations may be distributed among the one or more processors, not only residing within a single machine, but deployed across a number of machines. In some example embodiments, the processor or processors may be located in a single location (e.g., within a home environment, an office environment or as a server farm), while in other embodiments the processors may be distributed across a number of locations.

[0322] The one or more processors may also operate to support performance of the relevant operations in a "cloud computing" environment or as a "software as a service" (SaaS). For example, at least some of the operations

may be performed by a group of computers (as examples of machines including processors), these operations being accessible via a network (e.g., the Internet) and via one or more appropriate interfaces (e.g., Application Program Interfaces (APIs).)

ELECTRONIC APPARATUS AND SYSTEM

[0323] Example embodiments may be implemented in digital electronic circuitry, or in computer hardware, firmware, software, or in combinations of them. Example embodiments may be implemented using a computer program product, e.g., a computer program tangibly embodied in an information carrier, e.g., in a machine-readable medium for execution by, or to control the operation of, data processing apparatus, e.g., a programmable processor, a computer, or multiple computers.

[0324] A computer program can be written in any form of programming language, including compiled or interpreted languages, and it can be deployed in any form, including as a stand-alone program or as a module, subroutine, or other unit suitable for use in a computing environment. A computer program can be deployed to be executed on one computer or on multiple computers at one site or distributed across multiple sites and interconnected by a communication network.

[0325] In example embodiments, operations may be performed by one or more programmable processors executing a computer program to perform functions by operating on input data and generating output. Method operations can also be performed by, and apparatus of example embodiments may be implemented as, special purpose logic circuitry, e.g., a field programmable gate array (FPGA) or an application-specific integrated circuit (ASIC).

[0326] The computing system can 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. In embodiments deploying a programmable computing system, it will be appreciated that that both hardware and software architectures require consideration. Specifically, it will be appreciated that the choice of whether to implement certain functionality in permanently configured

hardware (e.g., an ASIC), in temporarily configured hardware (e.g., a combination of software and a programmable processor), or a combination of permanently and temporarily configured hardware may be a design choice. Below are set out hardware (e.g., machine) and software architectures that may be deployed, in various example embodiments.

EXAMPLE MACHINE ARCHITECTURE AND MACHINE-READABLE MEDIUM

[0327] Fig. 65 is a block diagram of machine in the example form of a computer system 6500 within which instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed. In alternative embodiments, the machine operates as a standalone device or may be connected (e.g., networked) to other machines. In a networked deployment, the machine may operate in the capacity of a server or a client machine in server-client network environment, or as a peer machine in a peer-to-peer (or distributed) network environment. The machine may be a personal computer (PC), a tablet PC, a set-top box (STB), a Personal Digital Assistant (PDA), a cellular telephone, a web appliance, a network router, switch or bridge, or any machine capable of executing instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while only a single machine is illustrated, the term "machine" shall also be taken to include any collection of machines that individually or jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein.

[0328] The example computer system 6500 includes a processor 6502 (e.g., a central processing unit (CPU), a graphics processing unit (GPU) or both), a main memory 6504 and a static memory 6506, which communicate with each other via a bus 6508. The computer system 6500 may further include a video display unit 6510 (e.g., a liquid crystal display (LCD) or a cathode ray tube (CRT)). The computer system 6500 also includes an alphanumeric input device 6512 (e.g., a keyboard or a touch-sensitive display screen), a user interface (UI) navigation device 6514 (e.g., a mouse), a disk drive unit 6516, a signal generation device 6518 (e.g., a speaker) and a network interface device 6520.

MACHINE-READABLE MEDIUM

[0329] The disk drive unit 6516 includes a machine-readable medium 6522 on which is stored one or more sets of instructions and data structures (e.g., software) 6524 embodying or utilized by any one or more of the methodologies or functions described herein. The instructions 6524 may also reside, completely or at least partially, within the main memory 6504 and/or within the processor 6502 during execution thereof by the computer system 6500, the main memory 6504 and the processor 6502 also constituting machine-readable media.

[0330] While the machine-readable medium 6522 is shown in an example embodiment to be a single medium, the term "machine-readable medium" may include a single medium or multiple media (e.g., a centralized or distributed database, and/or associated caches and servers) that store the one or more instructions or data structures. The term "machine-readable medium" shall also be taken to include any tangible medium that is capable of storing, encoding or carrying instructions for execution by the machine and that cause the machine to perform any one or more of the methodologies of the present invention, or that is capable of storing, encoding or carrying data structures utilized by or associated with such instructions. The term "machine-readable medium" shall accordingly be taken to include, but not be limited to, solid-state memories, and optical and magnetic media. Specific examples of machine-readable media include non-volatile memory, including by way of example semiconductor memory devices, e.g., Erasable Programmable Read-Only Memory (EPROM), Electrically Erasable Programmable Read-Only Memory (EEPROM), and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks.

TRANSMISSION MEDIUM

[0331] The instructions 6524 may further be transmitted or received over a communications network 6526 using a transmission medium. The instructions 6524 may be transmitted using the network interface device 6520 and any one of a number of well-known transfer protocols (e.g., HTTP). Examples of communication networks include a local area network ("LAN"), a wide area network ("WAN"), the Internet, mobile telephone networks, Plain Old Telephone

(POTS) networks, and wireless data networks (e.g., WiFi and WiMax networks). The term "transmission medium" shall be taken to include any intangible medium that is capable of storing, encoding or carrying instructions for execution by the machine, and includes digital or analog communications signals or other intangible media to facilitate communication of such software.

[0332] Although an embodiment has been described with reference to specific example embodiments, it will be evident that various modifications and changes may be made to these embodiments without departing from the broader spirit and scope of the invention. Accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense. The accompanying drawings that form a part hereof, show by way of illustration, and not of limitation, specific embodiments in which the subject matter may be practiced. The embodiments illustrated are described in sufficient detail to enable those skilled in the art to practice the teachings disclosed herein. Other embodiments may be utilized and derived therefrom, such that structural and logical substitutions and changes may be made without departing from the scope of this disclosure. This Detailed Description, therefore, is not to be taken in a limiting sense, and the scope of various embodiments is defined only by the appended claims, along with the full range of equivalents to which such claims are entitled.

[0333] Such embodiments of the inventive subject matter may be referred to herein, individually and/or collectively, by the term "invention" merely for convenience and without intending to voluntarily limit the scope of this application to any single invention or inventive concept if more than one is in fact disclosed. Thus, although specific embodiments have been illustrated and described herein, it should be appreciated that any arrangement calculated to achieve the same purpose may be substituted for the specific embodiments shown. This disclosure is intended to cover any and all adaptations or variations of various embodiments. Combinations of the above embodiments, and other embodiments not specifically described herein, will be apparent to those of skill in the art upon reviewing the above description.

CLAIMS

What is claimed is:

1. A method comprising:
generating a ratings matrix including matrix values, each row of the ratings matrix identifying one of a plurality of users, each column of the ratings matrix identifying one of a plurality of items, and each of the matrix values corresponding to a known affinity rating describing a degree of affinity associated with one of the users and one of the items, wherein the ratings matrix includes a missing entry representing an unknown affinity rating; and
generating, using one or more processors, a revised ratings matrix by factoring the ratings matrix into a user matrix and an item matrix, the revised ratings matrix being the product of the user matrix and the item matrix and including at least one entry representing a predicted affinity rating in place of the missing entry.
2. The method of claim 1, wherein each row of the user matrix includes user parameters describing each user in the ratings matrix, and wherein each column of the item matrix includes item parameters describing each item in the ratings matrix.
3. The method of claim 2, further comprising:
inserting a new row identifying a new user into the ratings matrix, the new row including a known affinity rating describing a degree of affinity associated with the new user; and
estimating user parameters describing the new user, based on the known affinity rating describing the degree of affinity associated with the new user and the item parameters included in the item matrix.
4. The method of claim 3, further comprising:
predicting an affinity rating describing a degree of affinity associated with the new user, based on the estimated user parameters and the item parameters included in the item matrix.

5. The method of claim 2, further comprising:

inserting a new column identifying a new item into the ratings matrix, the new column including a known affinity rating describing a degree of affinity associated with the new item; and

estimating item parameters describing the new item, based on the known affinity rating describing the degree of affinity associated with the new item and the user parameters included in the user matrix.

6. The method of claim 5, further comprising:

predicting an affinity rating describing a degree of affinity associated with the new item, based on the estimated item parameters and the user parameters included in the user matrix.

7. The method of claim 1, further comprising:

receiving a new affinity rating describing a degree of affinity associated with a particular user and a particular item; and

revising user parameters describing the particular user in the user matrix based on the new affinity rating.

8. The method of claim 1, further comprising:

receiving a new affinity rating describing a degree of affinity associated with a particular user and a particular item; and

revising item parameters describing the particular item in the item matrix based on the new affinity rating.

9. The method of claim 1, further comprising:

estimating a confidence value associated with at least a portion of the known affinity ratings in the ratings matrix; and

generating a confidence matrix that includes the confidence values.

10. The method of claim 9, further comprising:

estimating a probability that a specific user had an opportunity to rate a specific item, where the specific user and the specific item are associated with a missing affinity rating in the ratings matrix;

replacing the missing affinity rating in the ratings matrix with a low affinity rating; and

inserting a confidence value associated with the specific user and the specific item in the confidence matrix, where the confidence value has a value equal to the probability.

11. The method of claim 10, further comprising:

factoring the ratings matrix into the user matrix and the item matrix to generate the revised ratings matrix, by weighting the low affinity rating in the ratings matrix based on the confidence value in the confidence matrix.

12. The method of claim 9, further comprising:

estimating a probability that a missing affinity rating in the ratings matrix that describes an affinity associated with a specific user and a specific item corresponds to a negative affinity rating;

replacing the missing affinity rating in the ratings matrix with a low affinity rating; and

inserting a confidence value associated with the specific user and the specific item in the confidence matrix, where the confidence value has a value equal to the probability.

13. The method of claim 12, further comprising:

factoring the ratings matrix into the user matrix and the item matrix to generate the revised ratings matrix, by weighting the low affinity rating in the ratings matrix based on the confidence value in the confidence matrix.

14. The method of claim 2, further comprising:

determining that a first user is similar to a second user, based on the user parameters in the user matrix describing the first and second users; and

generating a recommendation of a particular item for the second user, based on a recommendation of the particular item for the first user.

15. The method of claim 2, further comprising
determining that a first item is similar to a second item, based on the item parameters in the item matrix describing the first and second items; and
generating a recommendation of the second item for a particular user, based on a recommendation of the first item for the particular user.

16. The method of claim 1, further comprising
counting a number of common affinity ratings between a first user and a second user, wherein each of the common affinity ratings describe a degree of affinity associated with a first item or a second item that is determined to be similar to the first item;
determining that the first user is similar to the second user, based on the number of common affinity ratings.

17. A non-transitory machine-readable storage medium having embodied thereon instructions executable by one or more machines to perform operations comprising:
generating a ratings matrix including matrix values, each row of the ratings matrix identifying one of a plurality of users, each column of the ratings matrix identifying one of a plurality of items, and each of the matrix values corresponding to a known affinity rating describing a degree of affinity associated with one of the users and one of the items, wherein the ratings matrix includes a missing entry representing an unknown affinity rating; and
generating, using one or more processors, a revised ratings matrix by factoring the ratings matrix into a user matrix and an item matrix, the revised ratings matrix being the product of the user matrix and the item matrix and including at least one entry representing a predicted affinity rating in place of the missing entry.

18. An apparatus comprising:
a matrix generation module configured to generate a ratings matrix including matrix values, each row of the ratings matrix identifying one of a plurality of users, each column of the ratings matrix identifying one of a plurality

of items, and each of the matrix values corresponding to a known affinity rating describing a degree of affinity associated with one of the users and one of the items, wherein the ratings matrix includes a missing entry representing an unknown affinity rating; and

a prediction module, implemented by one or more processors, configured to generate a revised ratings matrix by factoring the ratings matrix into a user matrix and an item matrix, the revised ratings matrix being the product of the user matrix and the item matrix and including at least one entry representing a predicted affinity rating in place of the missing entry.

19. The apparatus of claim 18, wherein the prediction module is further configured to:

insert a new row identifying a new user into the ratings matrix, the new row including a known affinity rating describing a degree of affinity associated with the new user; and

estimate user parameters describing the new user, based on the known affinity rating describing the degree of affinity associated with the new user and item parameters included in the item matrix.

20. The apparatus of claim 18, wherein the prediction module is further configured to:

insert a new column identifying a new item into the ratings matrix, the new column including a known affinity rating describing a degree of affinity associated with the new item; and

estimate item parameters describing the new item, based on the known affinity rating describing the degree of affinity associated with the new item and user parameters included in the user matrix.




FIG. 1




FIG. 2






FIG. 3








FIG. 6

7/62




FIG. 7

8/ 62




FIG. 8

810

REMEMBERED ANSWERS		PUBLIC / PRIVATE	
QUESTION: BOXERS OR BRIEFS? ARE YOU DRAWN TO: ARE YOU GENERALLY MORE: ARE YOUR THOUGHTS MOSTLY: WHICH IS MORE TRUE OF YOURSELF: ARE YOU USUALLY MORE INTERESTED IN: WHICH CELEBRITY DO YOU FIND MOST ATTRACTIVE: HOW DO YOU MANAGE YOUR TIME?	ANSWER: BOXERS CREDIBLE ASSUMPTIONS CONCEPTUAL IMAGINATIVE PUT THINGS OFF GENERAL THEORY PHOTO LEISURELY	REMOVE: ☐ ☐ ☐ ☐ ☐ ☐ ☐ ☐	

FIG. 8A

10/ 62

RATE ITEMS
TELL US IF AN ITEM WOULD BE A GOOD
DECISION FOR YOU

20%
THIS ACTIVITY

RATE Q&A
JUDGE THE QUALITY OF Q&A ASKED

10%
THIS ACTIVITY

QUALITY CONTROL ITEMS
VERIFY THAT SUBMITTED ITEMS HAVE
QUALITY PHOTOS AND DESCRIPTIONS

10%
THIS ACTIVITY

IDENTIFY DUPLICATE ITEMS
MARK PAIRS OF ITEMS AS DIFFERENT,
DUPLICATES, OR VERY SIMILAR

10%
THIS ACTIVITY

IDENTIFY DUPLICATE Q&A
MARK Q&A PAIRS AS DIFFERENT,
DUPLICATES, OR VERY SIMILAR

20%
THIS ACTIVITY

RATE PRO/CONS
MARK USER SUBMITTED PRO / CONS
AS HELPFUL OR UNHELPFUL

20%
THIS ACTIVITY

SELECT AN ACTIVITY IN CAMERAS

SWITCH TOPIC

FIG. 9

902




FIG. 10

1100 /
1102 /

THERE IS A 78% CHANCE YOU'LL AGREE WITH THIS DECISION:

78%

IMAGE

45%
DESCRIPTION

EXPERIMENTAL
IMAGE
DESCRIPTION

DESCRIPTION TITLE

DETAILED DESCRIPTION

DO YOU AGREE WITH THIS DECISION?
YES NO

IMAGE

VARIATIONS

SEE ALL

SUGGESTED TOPICS

TOPIC

TOPIC

TOPIC

TOPIC

TOPIC

1312

FIG. 11

PROS & CONS FROM PEOPLE LIKE YOU

1100 /

23 MINUTES AGO -- 288 PEOPLE FOUND THIS REVIEW HELPFUL

USER NAME

USER IMAGE

DEMOGRAPHIC

PROS

CONS

DID THIS HELP?

YES

NO

23 MINUTES AGO -- 288 PEOPLE FOUND THIS REVIEW HELPFUL

USER NAME

USER IMAGE

DEMOGRAPHIC

PROS

CONS

DID THIS HELP?

YES

NO

WRITE YOUR OWN PROS AND CONS

PROS

CONS

SUBMIT

FIG. 11A

1200

TOPIC: POWER MOWERS

49%

MOWER #1
IMAGE

DESCRIPTION TITLE
DETAILED DESCRIPTION
LISTED AT 'LOCATION' FOR 'PRICE'

DO YOU AGREE WITH THIS
DECISION?
YES NO

36%

MOWER #2
IMAGE

DESCRIPTION TITLE
DETAILED DESCRIPTION
LISTED AT 'LOCATION' FOR 'PRICE'

DO YOU AGREE WITH THIS
DECISION?
YES NO

30%

MOWER #3
IMAGE

DESCRIPTION TITLE
DETAILED DESCRIPTION
LISTED AT 'LOCATION' FOR 'PRICE'

DO YOU AGREE WITH THIS
DECISION?
YES NO

DISPLAYING FIRST 3 RESULTS

FIG. 12




FIG. 13

16/ 62




FIG. 14

17/ 62





FIG. 15




1700



HOBBIES AND PETS:

WHAT'S THE BEST NAME FOR MY NEW PUPPY?



IS THIS NAME FOR A FEMALE, OR MALE DOG.
☐ FEMALE
☐ MALE
 SKIP THIS QUESTION
 ADVANCED | EDIT

'USERNAME' CREATED THIS QUESTION ON 'DATE'

SKIP TO END

FIG. 17



HOBBIES AND PETS:

WHAT'S THE BEST NAME FOR MY NEW PUPPY?

1808

SELECTED TOPICS

/

WHAT'S THE BEST DOG BREED FOR ME?
98 PLAYS

IMAGE

WHICH DOG TOY WOULD MY POOCH ENJOY?
123 PLAYS

IMAGE

WHAT'S THE BEST DOG FOOD FOR MY DOG?
45 PLAYS

IMAGE

WHICH WALKING SHOE SHOULD I WEAR?
8 PLAYS

IMAGE

WHICH US SUMMER CAMP WOULD MY CHILD ENJOY?
33 PLAYS

IMAGE

#1

#2
1810

#3
1812

WILD CARD

ALL RESULTS



RUSTY

DESCRIPTION

1804

DO YOU LIKE THIS NAME?

☐ YES
☐ NO

1802

PROS AND CONS FROM PEOPLE LIKE YOU

PROS

USER IMAGE

RATING/
OPINION

CONS

USER IMAGE

RATING/
OPINION

HIGH SIMILARITY TO YOU

FIG. 18

1800

1900 /

EDIT RESULT TRAINING FOR 'RUSTY'

NEED HELP TRAINING THIS RESULT? READ OUR TRAINING FAQ.

RUSTY

DESCRIPTION

TRAINED QUESTIONS (5), IRRELEVANT QUESTIONS (0),
UNTRAINED QUESTIONS (1), LEARNED (3)

83% TRAINED

IS THIS NAME FOR A
FEMALE, OR MALE DOG?
☐ FEMALE
☒ MALE

HISTORY FIX THIS

HOW MANY SYLLABLES
DO YOU WANT TO HAVE?
☐ KEEP IT TO ONE
☒ NO MORE THAN 2
☒ MORE THAN 3 OK

HISTORY FIX THIS

HOW MANY POUNDS
SHOULD THE DOG BE?
☐ LESS THAN 10
☒ UP TO 50
☒ MORE THAN 50

HISTORY FIX THIS

THIS DOG IS
☒ SOCIAL
☐ SHY

HISTORY FIX THIS

THIS DOG IS
☒ CUTE
☐ STRONG

HISTORY FIX THIS

FIG. 19

21/ 62

2000 /

HOBBIES AND PETS

WHAT'S THE BEST NAME FOR MY NEW PUPPY?

#1

#2


#3

WILD CARD

ALL RESULTS



RUSTY



READ MORE URL

REVISIONS HISTORY

CANCEL

SAVE

SELECTED TOPICS

IMAGE

WHAT'S THE BEST DOG BREED FOR ME?
98 PLAYS

IMAGE

WHICH DOG TOY WOULD MY POOCH ENJOY?
123 PLAYS

IMAGE

WHAT'S THE BEST DOG FOOD FOR MY DOG?
45 PLAYS

IMAGE

WHICH WALKING SHOE SHOULD I WEAR?
8 PLAYS

IMAGE

WHICH US SUMMER CAMP WOULD MY CHILD ENJOY?
33 PLAYS

FIG. 20

2100

FIG. 21

VIEW VERSION OF "CANON REBEL T1i" (RESULT)

REVISION AS OF 12:28, 27 MAR 2009
(REVERT TO THIS ONE)
JON (DISCUSS | CONTRIBS)

RESULT NAME:
CANON REBEL T1i

RESULT IMAGE:
IMAGE

DESCRIPTION:
DESCRIPTION

READ MORE
URL
URL


REVISION AS OF 08:30, 14 APR 2009
STAFF (DISCUSS | CONTRIBS)

RESULT NAME:
CANON REBEL T1i

RESULT IMAGE:
IMAGE

DESCRIPTION:
DESCRIPTION


READ MORE
URL
URL




229

2300

FIG. 23





27/62

2500

TEACH US ABOUT YOU

DO YOU WEAR YOUR WATCH ON
YOUR RIGHT OR LEFT WRIST?

☐ DON'T WEAR A WATCH
☐ RIGHT
☐ LEFT

SKIP THIS QUESTION

FIG. 25

2600 /

FIG. 26

ACTIVITY FEED

IMAGE	'USERNAME' SAID YOUR PRO/CON FOR THE TOYOTA PRIUS WAS HELPFUL	40 MINUTES AGO
IMAGE	'USERNAME' ADDED 3 RESULTS IN 'WHICH MODERN FURNITURE SHOULD I BUY?'	AT 9:52 AM


YOU EARNED 1433 BANJOS IN WEEK ENDING 5/10/09 AT 12:07 AM

IMAGE	'USERNAME' ADDED 2 RESULT IN 2 TOPICS	YESTERDAY
IMAGE	'USERNAME' ADDED 2 RESULTS TO 3 QUESTIONS IN THE WORKSHOP	YESTERDAY
IMAGE	'USERNAME' ADDED 1 RESULT AND 2 PRO/CONS IN 3 TOPICS	YESTERDAY
IMAGE	THE TOPIC 'HOW TO GET ORGAINIZED', CREATED BY 'USERNAME' WAS PROMOTED OUT OF THE WORKSHOP	YESTERDAY

29/ 62

SEARCH: FIOS	2702	/	
RESULTS FOR: WHAT ISP SHOULD I USE? 1. VERIZON FIOS -- VERIZON FIOS, WHICH STANDS FOR FIBER OPTIC SERVICE, IS AN INTERNET, TELEPHONE, AND TV SERVICE THAT IS PRESENTLY OFFERED IN SOME AREAS OF THE U.S. BY VERIZON. VERIZON HAS ATTRACTED CONSUMER AND MEDIA ATTENTION IN THE AREA OF BROADBAND... 2. U-VERSE 3. COMCAST CABLE MODEM			
	2704	/	SEE ALL RESULTS
RESULTS FOR: WHAT US SATELLITE/CABLE SERVICE SHOULD I GET? 1. VERIZON FIOS -- VERIZON FIOS, IS A BUNDLED COMMUNICATIONS (INTERNET, TELEPHONE, AND TV) SERVICE, OPERATING OVER A FIBER-OPTIC COMMUNICATIONS NETWORK, THAT IS PRESENTLY OFFERED IN SOME AREAS OF THE U.S. BY VERIZON. VERIZON HAS ATTRACTED CONSUMER AND MEDIA ATTENTION IN THE AREA... 2. TIME WARNER CABLE 3. CHARTER COMMUNICATIONS			
			SEE ALL RESULTS

FIG. 27



WHAT NEW CAR SHOULD I BUY?

IMAGE

WHICH ONE OF YOUR ANSWERS
IS MOST IMPORTANT?


☐ UNDER \$18,000

☐ SUV

☐ PRACTICAL

SKIP THIS QUESTION

FIG. 29



CORRELATIONS FOR MARTHA STEWART'S FOLLOWERS				
IMAGE	COMPARED TO ALL TWITTER USERS, PEOPLE WHO FOLLOW MARTHA STEWART ARE: SOMEWHAT MORE LIKELY TO ANSWER....			
THESE ARE THE BATHROOM SIGNS AT THE SCIENC FICTION MUSEUM IN SEATTLE. WHICH DOOR WOULD YOU ENTER?	WITH	EVERY DAY	IMAGE	NOTHING. THAT WOULD SPOIL THE LINE IN MY CLOTHING. FEMALE
	WITH	YES		
	WITH			
	WITH			
	WITH			
HOW FREQUENTLY DO YOU WEAR MAKEUP?	WITH	WITHIN THE LAST YEAR	IMAGE	I'M WEARING VINTAGE RIGHT NOW! YES, I ROUTINELY MAKE...
	WITH			
	WITH			
	WITH			
	WITH			
DOES CELEBRITY GOSSIP INTEREST YOU?	WITH			
	WITH			
	WITH			
	WITH			
	WITH			
DO YOU CARRY ANYTHING IN YOUR POCKETS?	WITH			
	WITH			
	WITH			
	WITH			
	WITH			
ARE YOU MALE OR FEMALE?	WITH			
	WITH			
	WITH			
	WITH			
	WITH			
COMPARED TO ALL TWITTER USERS, PEOPLE WHO FOLLOW MARTHA STEWART ARE: SOMEWHAT MORE LIKELY TO ANSWER....	WITH			
	WITH			
	WITH			
	WITH			
	WITH			
WHICH WALLET IS CLOSEST TO YOURS?	WITH			
	WITH			
	WITH			
	WITH			
	WITH			
WHEN WAS THE LAST TIME YOU HAD A MANICURE?	WITH			
	WITH			
	WITH			
	WITH			
	WITH			
DO YOU EVER WEAR USED OR VINTAGE CLOTHES?	WITH			
	WITH			
	WITH			
	WITH			
	WITH			
SINCE BECOMING AN ADULT, HAVE YOU EVER HAND-CRAFTED A GIFT FOR SOMEONE?	WITH			
	WITH			
	WITH			
	WITH			
	WITH			
WHO DO YOU FIND THE MOST ATTRACTIVE?	WITH			
	WITH			
	WITH			
	WITH			
	WITH			

FIG. 31

34/ 82




FIG. 32

35/ 62




FIG. 33




FIG. 34

37/62




FIG. 35

38/ 62




FIG. 36

39/ 62




FIG. 37

40/62



FIG. 38

41/62



FIG. 39

42/ 62



FIG. 40




FIG. 41

44/ 82




FIG. 42

45/ 62





FIG. 43

46/62

**Fig. 44**

4500

	Smartphone	Tablet	Lemonade	Scooter
Users that like	George	George	Miguel	Miguel
Users that don't like	Miguel	Robert	Robert	George

4502

	George	Miguel	Robert
Items liked	Smartphone Tablet	Lemonade Scooter	-
Items not liked	Scooter	Smartphone	Tablet Lemonade

Fig. 45

4600

RATINGS MATRIX R (k x j matrix)

		j columns			
		j1 (Smartphone)	j2 (Tablet)	j3 (Lemonade)	j4 (Scooter)
k rows	k1 (George)	r11	r12	-	r14
	k2 (Robert)	-	r22	r23	-
	k3 (Miguel)	r31	-	r33	r34

Fig. 46




Fig. 47

50/62

5700

ITEM MATRIX I (p x j matrix)

j columns

		j1	j2	j3	j4
		(Smartphone)	(Tablet)	(Lemonade)	(Scooter)
p rows	ip1 (item parameter 1)	i11	i12	i13	i14
	ip2 (item parameter 2)	i21	i22	i23	i24
	ip3 (item parameter 3)	i31	i32	i33	i34
	ip4 (item parameter 4)	i41	i42	i43	i44
	ip5 (item parameter 5)	i51	i52	i53	i54

Fig. 48





Fig. 49

52/62

**Fig. 50**

53/62

Fig. 51

RATINGS MATRIX R ($k \times j$ matrix)
 j columns

5100

		j1 (Smartphone)	j2 (Tablet)	j3 (Lemonade)	j4 (Scooter)
k rows	k1 (George)	r_{11}	r_{12}	-	r_{14}
	k2 (Robert)	-	r_{22}	r_{23}	-
	k3 (Miguel)	r_{31}	-	r_{33}	r_{34}
	k4 (Sven)	r_{41}	r_{42}	-	r_{44}

Fig. 52

USER MATRIX U ($r \times p$ matrix)
 p columns

5200

		up1 (User Param. 1)	up2 (User Param. 2)	up3 (User Param. 3)	up4 (User Param. 4)	up5 (User Param. 5)
k rows	k1 (George)	u_{11}	u_{12}	u_{13}	u_{14}	u_{15}
	k2 (Robert)	u_{21}	u_{22}	u_{23}	u_{24}	u_{25}
	k3 (Miguel)	u_{31}	u_{32}	u_{33}	u_{34}	u_{35}
	k4 (Sven)	u_{41}	u_{42}	u_{43}	u_{44}	u_{45}

Fig. 53

54/62

RATINGS MATRIX R ($k \times j$ matrix)

j columns

	$j1$ (Smartphone)	$j2$ (Tablet)	$j3$ (Lemonade)	$j4$ (Scooter)	$j5$ (Printer)
$k1$ (George)	$r11$	$r12$	-	$r14$	$r15$
$k2$ (Robert)	-	$r22$	$r23$	-	-
$k3$ (Miguel)	$r31$	-	$r33$	$r34$	$r35$

k rows

5300

Fig. 54

ITEM MATRIX I ($p \times j$ matrix)



j columns

	$j1$ (Smartphone)	$j2$ (Tablet)	$j3$ (Lemonade)	$j4$ (Scooter)	$j5$ (Printer)
$ip1$ (item parameter 1)	$i11$	$i12$	$i13$	$i14$	$i15$
$ip2$ (item parameter 2)	$i21$	$i22$	$i23$	$i24$	$i25$
$ip3$ (item parameter 3)	$i31$	$i32$	$i33$	$i34$	$i35$
$ip4$ (item parameter 4)	$i41$	$i42$	$i43$	$i44$	$i45$
$ip5$ (item parameter 5)	$i51$	$i52$	$i53$	$i54$	$i55$



p rows

5400

55/62

Fig. 55**Fig. 56**

56/62


Fig. 57**Fig. 58**

5300


CONFIDENCE MATRIX C (k x j matrix)

		j columns			
		j1 (Smartphone)	j2 (Tablet)	j3 (Lemonade)	j4 (Scooter)
k rows	k1 (George)	c11	c12	-	c14
	k2 (Robert)	-	c22	c23	-
	k3 (Miguel)	c31	-	c33	c34



Fig. 59

**Fig. 60**


59/62

**Fig. 61**

60/62

Fig.62**Fig. 63**

61/62

**Fig. 64**




Fig.65

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US12/46536

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 17/30 (2012.01)

USPC - 707/748, 737

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; G06N 7/02 (2012.01)

USPC: 705/347; 707/748, 737, 780; 706/52

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

MicroPatent (US-G, US-A, EP-A, EP-B, WO, JP-bib, DE-C,B, DE-A, DE-T, DE-U, GB-A, FR-A); DialogPro (Derwent, INSPEC, NTIS, PASCAL, Current Contents Search, Dissertation Abstracts Online, Inside Conferences); IP.com; Google; Matrix, affinity, rating, correlate, user, item

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X — Y	US 2009/0299996 A1 (YU K et al.) December 3, 2009, paragraphs [0001], [0011], [0016]-[0018], [0092]	1, 2, 17 and 18 3-16, 19 and 20
Y	US 7685232 B2 (GIBBS S et al.) March 23, 2010, column 5, lines 56-63, column 7, lines 26-34, column 8, lines 46-52, column 9, lines 26-30, 40-62, column 11, lines 15-21, column 11, line 63 to column 2, line 18	3-8, 19 and 20
Y	US 2011/0106817 A1 (PAN R et al.) May 5, 2011, paragraphs [0020], [0022], [0030]	9-13
Y	US 2010/0268661 A1 (LEVY K et al.) October 21, 2010, paragraphs [0536], [0567], [0569]	14 and 15
Y	US 6092049 A (CHISLENKO A et al.) July 18, 2000, column 8, line 66 to column 9, line 11	16
A	US 7475027 B2 (BRAND M) January 6, 2009, entire document	1-20

☐ Further documents are listed in the continuation of Box C.

* 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

09 September 2012 (09.09.2012)

Date of mailing of the international search report

01 OCT 2012

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents

P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Shane Thomas

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774