



(51) International Patent Classification:

G06F 17/30 (2006.01) *G06T 7/00* (2006.01)
G06F 3/00 (2006.01) *G06F 15/16* (2006.01)

(21) International Application Number:

PCT/US2013/034470

(22) International Filing Date:

28 March 2013 (28.03.2013)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

13/434,738 29 March 2012 (29.03.2012) US

(71) Applicant (for all designated States except US): **YAHOO! INC.** [US/US]; 701 First Avenue, Sunnyvale, CA 94089 (US).

(72) Inventors; and

(71) Applicants (for US only): **SLANEY, Malcolm** [US/US]; 701 First Avenue, Sunnyvale, CA 94089 (US). **CHEN, Bee-chung** [—/US]; 701 First Avenue, Sunnyvale, CA 94089 (US). **AGARWAL, Deepak, K.** [IN/US]; 701 First Avenue, Sunnyvale, CA 94089 (US).

(74) Agent: **PITINGA, David, P.**; Martine Penilla Group, LLP, 710 Lakeway Drive, Suite 200, Sunnyvale, CA 94085 (US).


(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, 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, PA, 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).

Published:

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

(54) Title: FINDING ENGAGING MEDIA WITH INITIALIZED EXPLORE-EXPLOIT



(57) Abstract: Software for initialized explore-exploit creates a plurality of probability distributions. Each of these probability distributions is generated by inputting a quantitative description of one or more features associated with an image into a regression model that outputs a probability distribution for a measure of engagingness for the image. Each of the images is conceptually related to the other images. The software uses the plurality of probability distributions to initialize a multi-armed bandit model that outputs a serving scheme for each of the images. Then the software serves a plurality of the images on a web page displaying search results, based at least in part on the serving scheme.



FINDING ENGAGING MEDIA WITH INITIALIZED EXPLORE-EXPLOIT

BACKGROUND

[0001] As computing has become more mobile, client devices have shrunk in size, leaving little “real estate” on which to display text and images. This problem is particularly acute for images to be displayed on a faceted search engine result page (SERP), since most of the SERP is given over to textual search results. Thus, only a few images can be meaningfully shown in a SERP displayed on a client device such as a smartphone.

[0002] Ideally, these few images will include images that are relatively engaging, when compared to other images of the same subject. However, identifying engaging images in an automated fashion is difficult, even when the images are of celebrities and click-through data for the images is plentiful. When the subject of the images is less famous and click-through data for the images is sparse, human editors tend to be used instead of automation.

SUMMARY

[0003] In an example embodiment, a processor-executed method is described. According to the method, software for initialized explore-exploit creates a plurality of probability distributions. Each of these probability distributions is generated by inputting a quantitative description of one or more features associated with an image into a regression model that outputs a probability distribution for a measure of engagingness for the image. Each of the images is conceptually related to the other images. The software uses the plurality of probability distributions to initialize a multi-armed bandit model that outputs a serving scheme for each of the images. Then the software serves a plurality of the images on a web page displaying search results, based at least in part on the serving scheme.

[0004] In another example embodiment, an apparatus is described, namely, a computer-readable storage medium which persistently stores a program. The program might be a module in software for initialized explore-exploit. The program creates a plurality of probability distributions. Each of these probability distributions is generated by inputting a quantitative description of one or more features associated with an image into a regression model that outputs a probability distribution for a measure of engagingness for the image. Each of the images is conceptually related to the other images. The program uses the plurality of probability distributions to initialize a multi-armed bandit model that outputs a serving scheme for each of the images. Then the program serves a plurality of the images on a web page displaying search results, based at least in part on the serving scheme.

[0005] Another example embodiment also involves a processor-executed method. According to the method, software for initialized explore-exploit creates a plurality of probability distributions. Each of these probability distributions is generated by inputting a quantitative description of one or more features associated with a video into a regression model that outputs a probability distribution for a measure of engagingness for the video. Each of the videos is conceptually related. The software uses the plurality of probability distributions to initialize a multi-armed bandit model that outputs a serving scheme for each of the videos. Then the software serves a plurality of the videos on a web page displaying search results, based at least in part on the serving scheme.

[0006] Other aspects and advantages of the inventions will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, which illustrate by way of example the principles of the inventions.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0007] Figure 1 is a simplified network diagram that illustrates a website hosting a faceted search engine, in accordance with an example embodiment.
- [0008] Figure 2 is a software diagram that illustrates a functional modularization of software for finding engaging images, in accordance with an example embodiment.
- [0009] Figure 3 is a flowchart diagram that illustrates a process for generating engaging images for facets displayed on a search engine results page (SERP), in accordance with an example embodiment.
- [0010] Figure 4 shows a formula relating to a k-armed bandit model based on Bayesian inference, in accordance with an example embodiment.
- [0011] Figure 5 is a flowchart diagram that illustrates a process, using user feedback, for generating engaging images for facets displayed on a search engine results page (SERP), in accordance with an example embodiment.
- [0012] Figure 6 is a flowchart diagram that illustrates a process using user feedback to split a regression model, in accordance with an example embodiment.
- [0013] Figure 7 is a flowchart diagram that illustrates a process, using user feedback, for generating engaging audio samples for facets displayed on a search engine results page (SERP), in accordance with an example embodiment.
- [0014] Figure 8 is a flowchart diagram that illustrates a process, using user feedback, for generating engaging videos for facets displayed on a search engine results page (SERP), in accordance with an example embodiment.
- [0015] Figure 9 is a graph showing the relative performance of various multi-armed bandit models, in accordance with an example embodiment.

DETAILED DESCRIPTION

[0016] In the following description, numerous specific details are set forth in order to provide a thorough understanding of the exemplary embodiments. However, it will be apparent to one skilled in the art that the example embodiments may be practiced without some of these specific details. In other instances, process operations and implementation details have not been described in detail, if already well known.

[0017] One goal of the functionality described below is to find the most engaging image to go with any other piece of content (e.g., a web page, a description of a restaurant, or to illustrate a person), in the context of faceted search results. In an example embodiment, the functionality attempts to achieve this goal by looking for an image that produces relatively more mouse clicks than other related images. One might think of a mouse click as representing a request by the user for more information. Bad images will not reflect the user's interests, while good images will pique the user's interests. This is especially true for images used in conjunction with faceted search results, since only a small amount of screen “real estate” is used to summarize each search result.

[0018] Generally speaking, a relatively more engaging image of a person might show a face, while a relatively more engaging image of a restaurant might show a restaurant's ambiance. In an example embodiment, the functionality described below strives to learn the features of an engaging image through click-through data, which is a source of truth for the user's behavior. Click-through data associated with images tends to be more abundant for more popular search results, such as A-list celebrities and popular restaurants. More data makes for better models. But there are many search queries that will not have abundant click-through data associated with images. Thus, in an example embodiment, the functionality describe below trains an initial model to judge relatively

more engaging images of persons on data for A-list celebrities and then uses the trained model to predict relatively more engaging images for C-list celebrities.

[0019] In example embodiments, there might be enough available click-through data to train separate models for movie celebrities as opposed to sports celebrities, male movie celebrities as opposed to female movie celebrities, old female movie celebrities as opposed to young female movie celebrities, etc. In this regard, it will be appreciated that relatively more engaging images for sports celebrities who are curling athletes might look considerably different than relatively more engaging images of sports celebrities who are beach-volleyball athletes. On the other hand, there might not be much difference in more relatively engaging images of sports celebrities who are male hockey players and more relatively engaging images of sports celebrities who are female hockey players, since hockey players are often photographed in uniforms that only leave their faces exposed.

[0020] Figure 1 is a simplified network diagram that illustrates a website hosting a faceted search engine, in accordance with an example embodiment. As depicted in this figure, a personal computer 102 (e.g., a laptop or other mobile computer) and a mobile device 103 (e.g., a smartphone such as an iPhone, Android, Blackberry, etc., or a tablet computer such as an iPad, etc.) are connected by a network 101 (e.g., a wide area network (WAN) including the Internet, which might be wireless in part or in whole) with a website 104 hosting a search engine which supports faceted search, among other types of search. In an example embodiment, personal computer 102 and mobile device 103 might be used by end users who submit queries to the faceted search engine at website 104.

[0021] In an example embodiment, website 104 is composed of a number of servers connected by a network (e.g., a local area network (LAN) or a WAN) to each other in a cluster or

other distributed system which might run search-engine software, website software (e.g., web server software, database software, etc.), and distributed-computing and/or cloud software such as Map-Reduce, Google File System, Hadoop, Pig, CloudBase, etc. The servers are also connected (e.g., by a storage area network (SAN)) to persistent storage 105. Persistent storage 105 might include a redundant array of independent disks (RAID), in an example embodiment. Persistent storage 105 might be used to store data related to the search results, media (e.g., image files, audio files, and/or video files which might have been received from social-media websites), and models described in greater detail below.

[0022] Also connected to persistent storage 105 are the servers in cluster 106, which might run the initialized explore-exploit software described below. In an example embodiment, servers in cluster 106 are also connected through network 101 with website 107 which includes a cluster of servers that host social media (e.g., image files, audio files including container files for streaming, and/or video files including container files for streaming), such as images of celebrities, popular songs, and/or videos of newsworthy events. In an example embodiment, the cluster of servers in website 107 might also be connected (e.g., by a storage area network (SAN)) to persistent storage 108, which stores data related to social media (e.g., image files, audio files including container files for streaming, and/or video files including container files for streaming). Examples of website 107 include Flickr, YouTube, Vimeo, TwitPic, TwitVid, Y Frog image, Y Frog video, MySpace, Allmusic, Pandora, Spotify, etc. In an example embodiment, the servers in cluster 106 and in the cluster at website 107 might also run website software and distributed-computing and/or cloud software.

[0023] In an alternative example embodiment, the servers in website 104, cluster 106, website 107, and storage 105 might be hosted wholly or partially off-site in the cloud, e.g., as a platform-as-a-service (PaaS) or an infrastructure-as-a-service (IaaS).

[0024] Personal computers 102 and the servers in websites 104 and 107 and cluster 106 might include (1) hardware consisting of one or more microprocessors (e.g., from the x86 family or the PowerPC family), volatile storage (e.g., RAM), and persistent storage (e.g., a hard disk or solid-state drive), and (2) an operating system (e.g., Windows, Mac OS, Linux, Windows Server, Mac OS Server, etc.) that runs on the hardware. Similarly, in an example embodiment, mobile device 103 might include (1) hardware consisting of one or more microprocessors (e.g., from the ARM family), volatile storage (e.g., RAM), and persistent storage (e.g., flash memory such as microSD) and (2) an operating system (e.g., Symbian OS, RIM BlackBerry OS, iOS, webOS, Windows Mobile, Android, Linux, etc.) that runs on the hardware.

[0025] Also in an example embodiment, personal computer 102 and mobile device 103 might each include a browser as an application program or as part of an operating system. Examples of browsers that might execute on personal computer 102 include Internet Explorer, Mozilla Firefox, Safari, and Google Chrome. Examples of browsers that might execute on mobile device 103 include Safari, Mozilla Firefox, Android Browser, and webOS Browser. It will be appreciated that users of personal computer 102 and mobile device 103 might use browsers (e.g., through a graphical user interface or GUI) to communicate queries to and view search results from search engine software running on the servers at website 104. In an example embodiment, the search engine results might include the images, sound samples, and videos resulting from the software and processes described below. Examples of website 104 include Google, Bing, and Yahoo! Search, among others.

[0026] Figure 2 is a software diagram that illustrates a functional modularization of software for finding engaging images, in accordance with an example embodiment. In an example embodiment, the depicted functional modules might execute on the cluster of servers 106 shown in Figure 1. As depicted in Figure 2, initialization module 201 is an implementation (e.g., in R programming language) of a regression model (e.g., gradient-boosted decision trees) that has been trained using feature vectors (or other quantitative representations) from conceptually-related images with high click-through rates (e.g., images of the same A-list celebrity). The feature vector might result, for example, from applying a “bag of features” model, SIFT (Scale Invariant Feature Transform), SURF (Speeded Up Robust Feature), or a color histogram to an image which might have been received from a social-media website. Alternatively, the feature vector might include Boolean or enumerated values, e.g., a Boolean indicating whether an image is of a female or a male or whether the image includes cleavage or a man’s chest. In alternative example embodiments, the regression model might be based on (1) regression decision trees as implemented using Weka (Waikato Environment for Knowledge Analysis) or (2) support vector regression (SVR) machines, e.g., as implemented using the open-source LibSVM (Library for Support Vector Machines).

[0027] It will be appreciated that faceted classification and/or faceted search employs concepts (e.g., aspects, properties, or characteristics of a class or specific subject) to relate objects on multiple dimensions. For example, images of the same person are conceptually-related by the subject of the image. It will also be appreciated that click-through rate (CTR) is a measure of engagingness. And if images have a high CTR, there will be more data on which to base the regression model and consequently the regression model will tend to perform better.

[0028] Initialization module 201 receives feature vectors from conceptually-related images and maps them to corresponding probability distributions (e.g., with measures of central tendency

and variance) for expected CTRs. In an example embodiment, this might be done offline. Also, in an example embodiment, the conceptually-related images might be images of the same C-list celebrity. In alternative example embodiments, each probability distribution output by initialization model 201 might be a Gaussian distribution, a beta distribution, a gamma distribution, a Poisson distribution, etc.

[0029] In an example embodiment, the probability distributions output by initialization module 201 become inputs to online CTR module 202, which maintains the current probability distributions used as inputs for explore-exploit module 203. In an example embodiment, online CTR module might execute online, as suggested by its name. Online CTR module 202 passes the initial probability distributions as inputs to explore-exploit module 203, which might be implemented (e.g., using the R programming language) as a multi-armed (or k-armed) bandit model that employs, e.g., Bayesian inference with Lagrange relaxation. Such multi-armed bandit models are described in (1) co-owned U.S. Published Patent Application No. 2010/0125585 (U.S. Application No. 12/272,607) entitled “Conjoint Analysis with Bilinear Regression Models for Segmented Predictive Content Ranking” filed on November 17, 2008 (hereinafter “Conjoint Analysis”), and (2) the publication, Argawal et al., *Explore/Exploit Schemes for Web Content Optimization*, Ninth IEEE International Conference on Data Mining 2009 (ICDM '09) (hereinafter “Explore/Exploit Schemes”).

[0030] As depicted in Figure 2, the output of explore-exploit module 203 is a serving scheme for the conceptually-related images. According to the serving scheme (e.g., which might allocate a particular percentage of views to an image or display an image with a particular probability), some or all of the conceptually-related images will be displayed to users by serving engine 204 on web pages displaying faceted search results. In an example embodiment, serving

engine 204 gathers user feedback (e.g., mouse clicks) on the served images and inputs the user feedback to online CTR module 202, which incorporates the user feedback into the probability distributions for the served images (e.g., using one of the regression models described above), making them the current probability distributions used as inputs for explore-exploit module 203. In turn, these probability distributions are used to update the explore-exploit module 203, resulting in a revised serving scheme for the served images. In another example embodiment, the user feedback from the serving engine might be used to split initialization module 201 or the online CTR module 202 into multiple models, e.g., using image metadata.

[0031] It will be appreciated that a functional modularization similar to Figure 2 might be used to find engaging sound samples as discussed further below, in one alternative example embodiment. In another alternative example embodiment, a functional modularization similar to Figure 2 might be used to find engaging videos, as also discussed further below.

[0032] Figure 3 is a flowchart diagram that illustrates a process for generating engaging images for facets displayed on a search engine results page (SERP), in accordance with an example embodiment. In an example embodiment, this process might be performed by the initialized explore-exploit software executing on server cluster 106 as shown in Figure 1. As depicted in Figure 3, the software generates a quantitative description (e.g., feature vector) for each image in a group of conceptually-related images (e.g., images of same person), in operation 301. As indicated above, the conceptually-related images might have been collected from one or more social-media websites, e.g., using a search engine or a program such as mediaboxAdvanced. In operation 302, the software inputs each quantitative description into a regression model (e.g., gradient-boosted decision trees) trained using quantitative descriptions of images having a relatively large amount of data with respect to a measure of engagingness (e.g., click-through data on images of same

celebrity). As an output of the regression model, the software receives a probability distribution (e.g., that includes a measure of central tendency and a measure of variance) for each conceptually-related image, in operation 303. In operation 304, the software uses the probability distributions to initialize a multi-armed bandit model (e.g., an explore-exploit model based on Bayesian inference with Lagrange relaxation) that outputs a serving scheme for the group of conceptually-related images. And in operation 305, the software uses the serving scheme to serve some or all of the conceptually-related images (e.g., as thumbnail images) on a webpage displaying search results (e.g., on a search engine results page for faceted search).

[0033] Figure 4 shows a formula relating to a k-armed bandit model based on Bayesian inference, in accordance with an example embodiment. The formula 401 comes from Explore/Exploit Schemes and assumes a k-armed bandit where k is 2 (e.g., there are two items displayed to viewers) and where there are two intervals (e.g., time 0 and time 1). As explained in Explore/Exploit Schemes, the formula is an equivalence; that is to say, both sides of the formula 401 represent the expected total number of clicks in the two intervals. It will be appreciated that an optimal solution to this formula will find a decision x in the interval $[0,1]$ and a decision x_1 in the set of decisions X_1 (where x_1 and X_1 are functions of a random variable c) that maximizes the expected total number of clicks in the two intervals. A decision is a fraction of user visits allocated to a certain item, as noted in the annotation to formula 401 in Figure 4. It will be appreciated that the decisions x and x_1 constitute a “serving scheme”, as that term is used elsewhere in this disclosure. When k is greater than 2 (e.g., there are more than 2 items to be displayed to viewers), Lagrange relaxation simplifies the computation of the optimal solution, as also pointed out in Explore/Exploit Schemes.

[0034] Figure 5 is a flowchart diagram that illustrates a process, using user feedback, for generating engaging images for facets displayed on a search engine results page (SERP), in accordance with an example embodiment. In an example embodiment, this process might be performed by the initialized explore-exploit software executing on server cluster 106 as shown in Figure 1. As depicted in Figure 5, the software generates a quantitative description (e.g., feature vector) for each image in a group of conceptually-related images (e.g., images of same person), in operation 501. In operation 502, the software inputs each quantitative description into a regression model (e.g., gradient-boosted decision trees) trained using quantitative descriptions of images having a relatively large amount of data with respect to a measure of engagingness (e.g., click-through data on images of same celebrity). As an output of the regression model, the software receives a probability distribution (e.g., that includes a measure of central tendency and a measure of variance) for each conceptually-related image, in operation 503. In operation 504, the software uses the probability distributions to initialize a multi-armed bandit model (e.g., an explore-exploit model based on Bayesian inference with Lagrange relaxation) that outputs a serving scheme for the group of conceptually-related images. And in operation 505, the software uses the serving scheme to serve some or all of the conceptually-related images (e.g., as thumbnail images) on a webpage displaying search results (e.g., on a search engine results page for faceted search).

[0035] In operation 506, the software monitors user feedback (e.g., mouse clicks) on the served images and, in operation 507, the software incorporates user feedback (e.g., mouse clicks) into a probability distribution for each served image, e.g., at online CTR module 202 in Figure 2. In an alternative example embodiment, the user feedback might be a social signal such as a thumbs-up, like, or plus-one on the SERP or a share (or other post/broadcast such as a tweet) to a stream associated with a profile on social-networking website. In an example embodiment, the software

might use a regression model similar to that used in operation 502 to incorporate the user feedback into the probability distributions. In an alternative example embodiment, the software might use a different regression model, e.g., one of the other regression models discussed above, to incorporate the user feedback into the probability distributions. Then in operation 508, the software updates the multi-armed bandit model (e.g., an explore-exploit model based on Bayesian inference with Lagrange relaxation) with the probability distributions from the served images, for use in generating a revised serving scheme for the served images.

[0036] Figure 6 is a flowchart diagram that illustrates a process for using user feedback to split a regression model, in accordance with an example embodiment. In an example embodiment, this process might be performed by the initialized explore-exploit software executing on server cluster 106 as shown in Figure 1. As depicted in Figure 6, the software generates a quantitative description (e.g., feature vector) for each image in a group of conceptually-related images (e.g., images of same person), in operation 601. In operation 602, the software inputs each quantitative description into a regression model (e.g., gradient-boosted decision trees) trained using quantitative descriptions of images having a relatively large amount of data with respect to a measure of engagingness (e.g., click-through data on images of same celebrity). As an output of the regression model, the software receives a probability distribution (e.g., that includes a measure of central tendency and a measure of variance) for each conceptually-related image, in operation 603. In operation 604, the software uses the probability distributions to initialize a multi-armed bandit model (e.g., an explore-exploit model based on Bayesian inference with Lagrange relaxation) that outputs a serving scheme for the group of conceptually-related images. And in operation 605, the software uses the serving scheme to serve some or all of the conceptually-related images (e.g., as

thumbnail images) on a webpage displaying search results (e.g., on a search engine results page for faceted search).

[0037] In operation 606, the software monitors user feedback (e.g., mouse clicks) on the served images. In operation 607, the software incorporates the user feedback (e.g., mouse clicks) into a probability distribution for each served image, e.g., at online CTR module 202 in Figure 2. In an alternative example embodiment, the user feedback might be a social signal such as a thumbs-up, like, or plus-one on the SERP or a share (or other post/broadcast such as a tweet) to a stream associated with a profile on social-networking website. In an example embodiment, the software might use a regression model similar to that used in operation 602 to incorporate the user feedback into the probability distributions. In an alternative example embodiment, the software might use a different regression model, e.g., one of the other regression models discussed above, to incorporate the user feedback into the probability distributions. Then in operation 608, the software splits the regression model into multiple regression models, using metadata associated with the served images (e.g., descriptive information including keywords) and the user feedback from the served images. For example, if the metadata indicates that the subject of the served images is a male, the regression model might be split into a regression model for images of males and a regression model for images of females. The multiple regression models might then be used to generate initializing probability distributions from feature vectors extracted from a subsequent group of conceptually-related images.

[0038] It will be appreciated that this splitting of the regression model might be applied to initialization module 201 and/or online CTR module 202, in an example embodiment. An alternative example embodiment might provide for another form of “model splitting”, e.g., using Bayesian smoothing or shrinkage estimation as described by Pandey and Argawal et al., *Bandits for*

taxonomies: A model based approach, In Proceedings of the SIAM International Conference on Data Mining (2007), which is hereby incorporated by reference. Thus, given a pre-defined taxonomy of images (e.g., male and female), the software might use user feedback classified according to that taxonomy to “split” the explore-exploit model into multiple explore-exploit models.

[0039] Figure 7 is a flowchart diagram that illustrates a process, using user feedback, for generating engaging audio samples for facets displayed on a search engine results page (SERP), in accordance with an example embodiment. In an example embodiment, this process might be performed by the initialized explore-exploit software executing on server cluster 106 as shown in Figure 1. As depicted in Figure 7, the software generates a quantitative description (e.g., a feature vector) for each image in a group of conceptually-related audio samples (e.g., audio samples from the same song), in operation 701. The feature vector might include, for example, MFCCs (Mel-frequency cepstral coefficients) extracted from audio files received from a social-media website. In operation 702, the software inputs each quantitative description into a regression model (e.g., gradient-boosted decision trees) trained using quantitative descriptions of audio samples having a relatively large amount of data with respect to a measure of engagingness (e.g., play durations of audio samples from the same popular song). As an output of the regression model, the software receives a probability distribution (e.g., that includes a measure of central tendency and a measure of variance) for each conceptually-related audio sample, in operation 703. In operation 704, the software uses the probability distributions to initialize a multi-armed bandit model (e.g., an explore-exploit model based on Bayesian inference with Lagrange relaxation) that outputs a serving scheme for the group of conceptually-related audio samples. And in operation 705, the software uses the

serving scheme to serve some or all of the conceptually-related audio samples on a webpage displaying search results (e.g., on a search engine results page for faceted search).

[0040] In operation 706, the software monitors user feedback (e.g., mouse clicks) on the served audio samples and, in operation 707, the software incorporates user feedback (e.g., mouse clicks) into a probability distribution for each served audio sample, e.g., at online CTR module 202 in Figure 2. In an alternative example embodiment, the user feedback might be a social signal such as a thumbs-up, like, or plus-one on the SERP or a share (or other post/broadcast such as a tweet) to a stream associated with a profile on social-networking website. In an example embodiment, the software might use a regression model similar to that used in operation 702 to incorporate the user feedback into the probability distributions. In an alternative example embodiment, the software might use a different regression model, e.g., one of the other regression models discussed above, to incorporate the user feedback into the probability distributions. Then in operation 708, the software updates the multi-armed bandit model (e.g., an explore-exploit model based on Bayesian inference with Lagrange relaxation) with the probability distributions from the served audio samples, for use in generating a revised serving scheme for the served audio samples.

[0041] It will be appreciated that the process described in Figure 7 might be used to identify a hook in a song. Generally speaking, a hook is a musical idea (e.g., a short riff, passage, or phrase) which is used in popular music to make a song "catch the ear of the listener". In an alternative example embodiment, some of the operations in the process (e.g., operations 701-704) might be used in software that facilitates song composition rather than in software that generates audio samples for a faceted SERP. Likewise, some of the operations in the process described in Figure 5 (e.g., operations 501-504) might be used in "profile-picture" software rather than in software that generates images for a faceted SERP.

[0042] Figure 8 is a flowchart diagram that illustrates a process, using user feedback, for generating engaging videos for facets displayed on a search engine results page (SERP), in accordance with an example embodiment. In an example embodiment, this process might be performed by the initialized explore-exploit software executing on server cluster 106 as shown in Figure 1. As depicted in Figure 8, the software generates a quantitative description (e.g., a feature vector) for each video in a group of conceptually-related videos (e.g., videos of the same event or from the same film), in operation 801. The feature vector might include, for example, image and audio features extracted from container files received from a social-media website. In operation 802, the software inputs each quantitative description into a regression model (e.g., gradient-boosted decision trees) trained using quantitative descriptions of videos having a relatively large amount of data with respect to a measure of engagingness (e.g., play durations of videos of the same popular event or from the same popular film). As an output of the regression model, the software receives a probability distribution (e.g., that includes a measure of central tendency and a measure of variance) for each conceptually-related video, in operation 803. In operation 804, the software uses the probability distributions to initialize a multi-armed bandit model (e.g., an explore-exploit model based on Bayesian inference with Lagrange relaxation) that outputs a serving scheme for the group of conceptually-related videos. And in operation 805, the software uses the serving scheme to serve some or all of the conceptually-related videos on a webpage displaying search results (e.g., on a search engine results page for faceted search).

[0043] In operation 806, the software monitors user feedback (e.g., mouse clicks) on the served audio samples and, in operation 807, the software incorporates user feedback (e.g., mouse clicks) into a probability distribution for each served video, e.g., at online CTR module 202 in Figure 2. In an alternative example embodiment, the user feedback might be a social signal such as

a thumbs-up, like, or plus-one on the SERP or a share (or other post/broadcast such as a tweet) to a stream associated with a profile on social-networking website. In an example embodiment, the software might use a regression model similar to that used in operation 802 to incorporate the user feedback into the probability distributions. In an alternative example embodiment, the software might use a different regression model, e.g., one of the other regression models discussed above, to incorporate the user feedback into the probability distributions. Then in operation 808, the software updates the multi-armed bandit model (e.g., an explore-exploit model based on Bayesian inference with Lagrange relaxation) with the probability distributions from the served videos, for use in generating a revised serving scheme for the served videos.

[0044] Here again, it will be appreciated that some of the operations in the process described in Figure 8 (e.g., operations 801-804) might be used in software for composing film trailers rather than in software that generates videos for a faceted SERP.

[0045] Figure 9 is a graph depicting the relative performance of three images input to the process described in Figure 5, in accordance with an example embodiment. The three images were of the same A-list celebrity. In graph 901, the horizontal axis represents time in days and the vertical axis represents estimated CTR (e.g., as output by the online CTR module 202 and received as an input by the explore-exploit module 203 in Figure 2). Also, in graph 901, the circles represent estimated CTRs and the dashes represent their corresponding estimated standard deviations. At time 0, all three images have the same estimated CTR, namely, approximately 0.013. However, as days pass, the estimated CTR for the image represented by the topmost series of circles converges on that value, whereas the estimated CTR for the image represented by the bottommost series of circles is so low (approximately 0.009) that it is no longer shown to viewers after a few days. Similarly, after approximately 10 days, the estimated CTR for the image represented by the middle series of circles

is relatively poor, at approximately 0.010, and it, too, is no longer shown to viewers. It will be appreciated that as the images are shown to more and more viewers as the days pass, the click-through data for the images grows. And consequently, the estimated standard deviations of the estimated CTRs decline over time; that is to say, the estimates improve.

[0046] With the above embodiments in mind, it should be understood that the inventions might employ various computer-implemented operations involving data stored in computer systems. These operations are those requiring physical manipulation of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. Further, the manipulations performed are often referred to in terms, such as producing, identifying, determining, or comparing.

[0047] Any of the operations described herein that form part of the inventions are useful machine operations. The inventions also relate to a device or an apparatus for performing these operations. The apparatus may be specially constructed for the required purposes, such as the carrier network discussed above, or it may be a general purpose computer selectively activated or configured by a computer program stored in the computer. In particular, various general purpose machines may be used with computer programs written in accordance with the teachings herein, or it may be more convenient to construct a more specialized apparatus to perform the required operations.

[0048] The inventions can also be embodied as computer readable code on a computer readable medium. The computer readable medium is any data storage device that can store data, which can thereafter be read by a computer system. Examples of the computer readable medium include hard drives, network attached storage (NAS), read-only memory, random-access memory, CD-ROMs, CD-Rs, CD-RWs, DVDs, Flash, magnetic tapes, and other optical and non-optical data

storage devices. The computer readable medium can also be distributed over a network coupled computer systems so that the computer readable code is stored and executed in a distributed fashion.

[0049] Although example embodiments of the inventions have been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications can be practiced within the scope of the following claims. For example, some or all of the processes described above might be used with software for reinforcement learning that is not based on explore-exploit or multi-arm bandit models. Moreover, the operations described above can be ordered, modularized, and/or distributed in any suitable way. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the inventions are not to be limited to the details given herein, but may be modified within the scope and equivalents of the following claims. In the following claims, elements and/or steps do not imply any particular order of operation, unless explicitly stated in the claims or implicitly required by the disclosure.

What is claimed is:

CLAIMS

1. A method, comprising operations of:

creating a first plurality of probability distributions, wherein each of these probability distributions is generated by inputting a quantitative description of one or more features associated with a image into a regression model that outputs a probability distribution for a measure of engagingness for the image and wherein each of the images is conceptually related to the other images;

using the first plurality of probability distributions to initialize a multi-armed bandit model that outputs a serving scheme for each of the images; and

serving a plurality of the images on a web page displaying search results, based at least in part on the serving scheme, wherein each of the operations is executed by one or more processors.

2. The method of claim 1, further comprising operations of:

monitoring user feedback on the plurality of served images; and

creating a second plurality of probability distributions based at least in part on the user feedback; and

updating the multi-armed bandit model with the second plurality of probability distributions.

3. The method of claim 1, further comprising operations of:

monitoring user feedback on each of the plurality of served images; and

splitting the regression model into two or more regression models, based at least in part on the user feedback for the plurality of served images and image metadata associated with each of the plurality of served images.

4. The method of claim 2, wherein the user feedback is a mouse click.
5. The method of claim 1, wherein the measure of engagingness is click-through rate (CTR).
6. The method of claim 1, wherein the multi-armed bandit model employs Bayesian inference with Lagrange relaxation.
7. The method of claim 1, wherein the regression model is trained using images of celebrities.
8. The method of claim 1, wherein the regression model is based on gradient-boosted decision trees.
9. The method of claim 1, wherein the quantitative representation is a feature vector.
10. A computer-readable storage medium persistently storing a program, wherein the program, when executed, instructs a processor to perform the following operations:
 - create a first plurality of probability distributions, wherein each of these probability distributions is generated by inputting a quantitative description of one or more features associated with an image into a regression model that outputs a probability distribution for a measure of engagingness for the image and wherein each of the images is conceptually related to the other images;
 - use the first plurality of probability distributions to initialize a multi-armed bandit model that outputs a serving scheme for each of the images; and

serve a plurality of the images on a web page displaying search results, based at least in part on the serving scheme

11. The computer-readable storage medium of claim 10, further comprising operations of:

monitor user feedback on the plurality of served images; and

create a second plurality of probability distributions based at least in part on the user feedback; and

update the multi-armed bandit model with the second plurality of probability distributions.

12. The computer-readable storage medium of claim 10, further comprising operations of:

monitor user feedback on each of the plurality of served images; and

split the regression model into two or more regression models, based at least in part on the user feedback for the plurality of served images and image metadata associated with each of the plurality of served images.

13. The computer-readable storage medium of claim 11, wherein the user feedback is a mouse click.

14. The computer-readable storage medium of claim 10, wherein the measure of engagingness is click-through rate (CTR).

15. The computer-readable storage medium of claim 10, wherein the multi-armed bandit model employs Bayesian inference with Lagrange relaxation.

16. The computer-readable storage medium of claim 10, wherein the regression model is trained using images of celebrities.

17. The computer-readable storage medium of claim 10, wherein the regression model is based on gradient-boosted decision trees.

18. The computer-readable storage medium of claim 10, wherein the quantitative representation is a feature vector.

19. A method, comprising the operations of:

creating a first plurality of probability distributions, wherein each of these probability distributions is generated by inputting a quantitative description of one or more features associated with a video into a regression model that outputs a probability distribution for a measure of engagingness for the video and wherein each of the videos are conceptually related;

using the first plurality of probability distributions to initialize a multi-armed bandit model that outputs a serving scheme for each of the videos; and

serving a plurality of the videos on a web page displaying search results, based at least in part on the serving scheme, wherein each of the operations is executed by one or more processors.

20. The method of claim 19, further comprising operations of:

monitoring user feedback on the plurality of served videos; and

creating a second plurality of probability distributions based at least in part on the user feedback; and

updating the multi-armed bandit model with the second plurality of probability distributions.




Figure 1




Figure 2

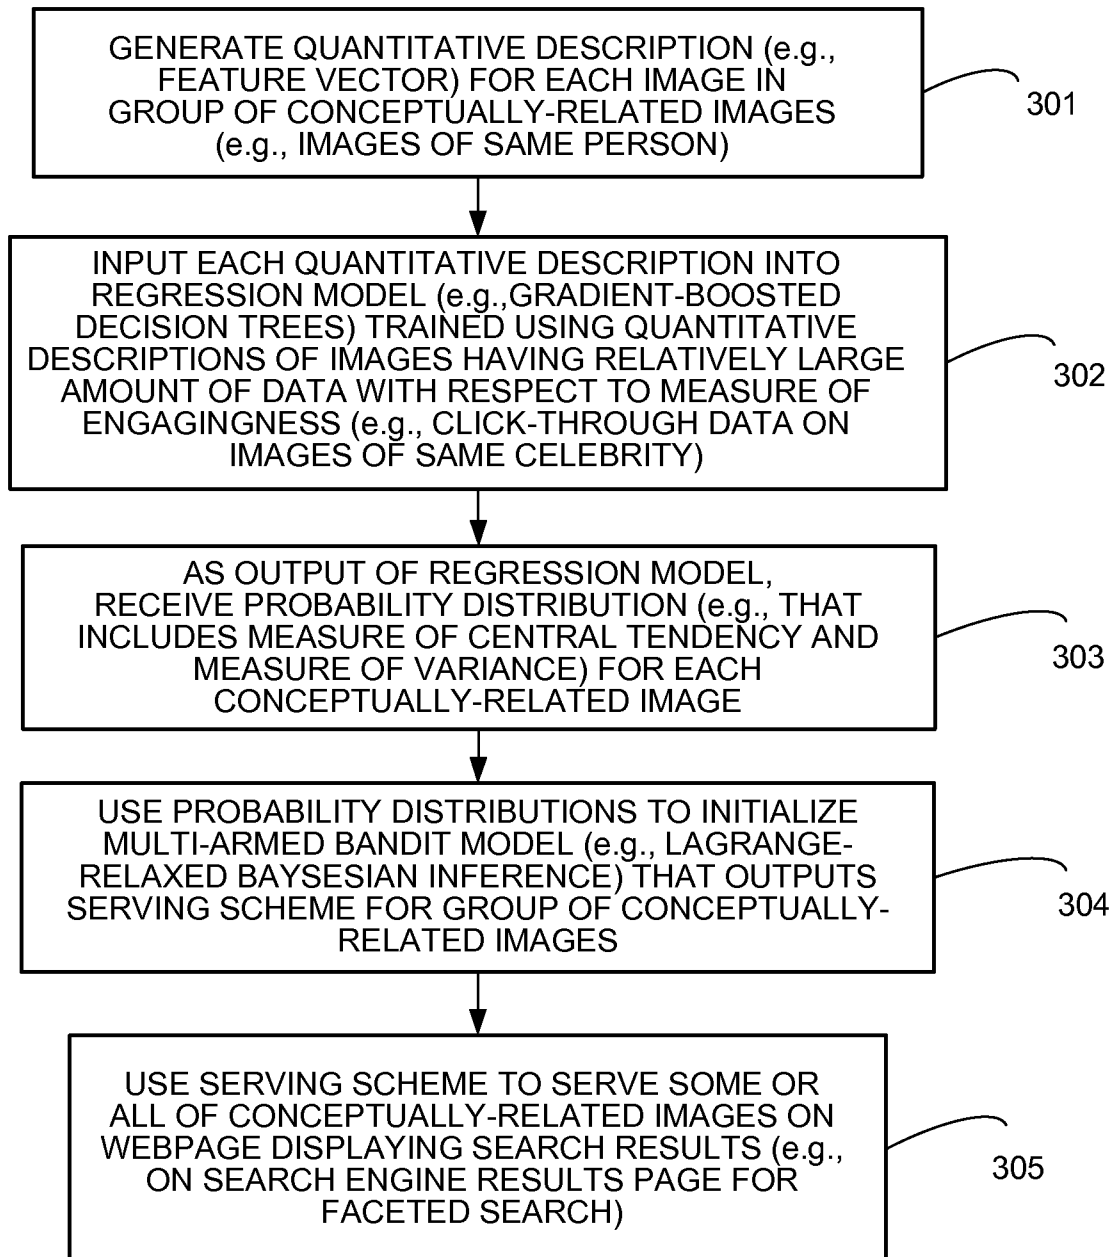


Figure 3

$$E[N_0(xp_0 + (1-x)q_0) + N_1(x_1p_1 + (1-x_1)q_1)]$$

$$= E[N_0x(p_0 - q_0) + N_1x_1(p_1 - q_1)] + q_0N_0 + q_1N_1$$

401

- N_0, N_1 : number of user visits at time 0 and 1.
- q_0, q_1 : CTR of the *certain* item at time 0 and 1.
- $p_0 \sim \mathcal{P}(\theta_0), p_1 \sim \mathcal{P}(\theta_1)$ are CTR distributions of the *uncertain* item at time 0 and 1 respectively.
- x, x_1 are the fractions of user visits allocated to the uncertain item at time 0 and 1; $(1-x)$ and $(1-x_1)$ are the fractions allocated to the certain item.
- c , a random variable, denotes the number of clicks that the uncertain item gets at time 0.

Figure 4




Figure 5




Figure 6




Figure 7




Figure 8




Figure 9

A. CLASSIFICATION OF SUBJECT MATTER**G06F 17/30(2006.01)i, G06F 3/00(2006.01)i, G06T 7/00(2006.01)i, G06F 15/16(2006.01)i**

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)

G06F 17/30; H04N 7/10; G06F 15/16; G06G 7/48; G06Q 30/00; G06N 5/02; G06F 15/18; G06F 3/00; G06T 7/00

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: probability distribution, regression model, multi-armed bandit model, explore, exploit, click through rate, feature vector, initialize

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2010-0191689 A1 (CORINNA CORTES et al.) 29 July 2010 See paragraphs [0003]-[0012], [0019], [0028]-[0034], [0036]-[0041], [0054]; and claims 1, 7, 12.	1-2, 4-11, 13-20
A		3, 12
Y	US 2010-0121624 A1 (H. SCOTT ROY et al.) 13 May 2010 See paragraphs [0002]-[0008], [0016]-[0018], [0025]-[0033], [0042]-[0043], [0059]-[0065]; and claims 8-9, 14.	1-2, 4-11, 13-20
Y	FILIP RADLINSKI et al. `Learning Diverse Ranking with Multi-Armed Bandits`. In: Proceedings of the 25th International Conference on Machine Learning, Helsinki, Finland, 2008. See pages 1, 3-5.	1-2, 4-11, 13-20
A	DEEPAK AGARWAL et al. `Explore/Exploit Schemes for Web Content Optimizatøn`. In: 9th IEEE International Conference on Data Mining 2009 (ICDM `09), December 2009, ISSN 1550-4786. See pages 1-5, 8.	1-20
A	US 2010-0241597 A1 (BEE-CHUNG CHEN et al.) 23 September 2010 See paragraphs [0001]-[0006], [0018]-[0028], [0037]-[0042], [0056]-[0060]; and claims 1-2.	1-20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of 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

26 June 2013 (26.06.2013)

Date of mailing of the international search report

28 June 2013 (28.06.2013)

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City,
302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

BYUN, Sung Cheal

Telephone No. 82-42-481-8262



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2013/034470

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 2010-026297 A1 (XTRACT OY) 11 March 2010 See page 1, line 5 - page 11, line 20; page 13, lines 21-31; page 14, line 6 - page 15, line 6; page 15, lines 21-24; and claims 1, 15.	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/034470

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010-0191689 A1	29.07.2010	EP 2382782 A1 EP 2382782 A4 US 2012-272259 A1 US 8301498 B1 WO 2010-087909 A1	02.11.2011 01.05.2013 25.10.2012 30.10.2012 05.08.2010
US 2010-0121624 A1	13.05.2010	US 2012-303349 A1 US 8244517 B2	29.11.2012 14.08.2012
US 2010-0241597 A1	23.09.2010	None	
WO 2010-026297 A1	11.03.2010	FI 20085843 A0 FI 20085843 D0 FI 20095167 A0 FI 20095167 D0	08.09.2008 08.09.2008 20.02.2009 20.02.2009