

(51) International Patent Classification:
G06F 17/30 (2006.01)(21) International Application Number:
PCT/US2009/061713(22) International Filing Date:
22 October 2009 (22.10.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/257,991 24 October 2008 (24.10.2008) US(71) Applicant (for all designated States except US): **YAHOO! INC.** [US/US]; 701 First Avenue, Sunnyvale, California 94089 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **MURDOCK, Vanessa** [US/ES]; c. Villar 57 3-2, E-08041 Barcelona (ES). **VAN ZWOL, Roelof** [NL/ES]; Vaixell Ma Assumpta, E-08912 Barcelona (ES). **PUEYO, Lluís Garcia** [ES/ES]; c/ Benet Corada 59, 2n 3a, E-08173 Sant Cugatdel Vallas (ES). **RAMIREZ CAMPS, Georgina** [ES/ES]; Ocata 1, E-08003 Barcelona (ES).(74) Agents: **CHEE, Robert, S.** et al.; 2055 Gateway Place, Ste. 550, San Jose, California 95110 (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, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,

[Continued on next page]

(54) Title: METHODS FOR IMPROVING THE DIVERSITY OF IMAGE SEARCH RESULTS

FIG. 2



(57) Abstract: Techniques are described to increase the diversity or focus of image search results. A user submits an original query to search for images. A server generates a first results set by executing the original query using metadata associated with each image. The server selects, from the first results set, a specified number of results ranked highest and generates a list of terms from the metadata of each of the results selected. The terms may be only the tags of the results. The server generates an updated query using terms in the list that may be weighted based on the frequency of the term in the list or include only a specified number of the highest occurring terms in the list. The server generates a second results set by executing the updated query using metadata associated with each image. The second results set is then stored and displayed to the user.



MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, **Published:**

TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,

ML, MR, NE, SN, TD, TG).

— *without international search report and to be republished
upon receipt of that report (Rule 48.2(g))*

METHODS FOR IMPROVING THE DIVERSITY OF IMAGE SEARCH RESULTS

FIELD OF THE INVENTION

[0001] The present invention relates to image searching.

BACKGROUND

[0002] The approaches described in this section are approaches that could be pursued, but not necessarily approaches that have been previously conceived or pursued. Therefore, unless otherwise indicated, it should not be assumed that any of the approaches described in this section qualify as prior art merely by virtue of their inclusion in this section.

[0003] Search engines help a user to locate information. Using a search engine, a user may enter one or more search query terms and obtain a list of resources that contain or are associated with subject matter that matches those search query terms. To find the most relevant files, search engines typically attempt to select, from among a plurality of files, files that include many or all of the words that a user entered into a search request. This works well when the files searched contain a large quantity of text where relevance and context may be determined. Images, however, present problems because often the only way to determine the content and context of an image is through small snippets of text associated with the image.

[0004] As photo sharing services on the Internet have become more popular, obtaining more effective search results of those images has become increasingly important. Photo sharing services allow users to upload and share digital images with family, friends, and the general public.

[0005] Users may be given the opportunity to provide annotations to each particular photo uploaded. These annotations, which may also be referred to as user-generated content, may define a title, description, and a set of tags for the photo. The set of tags might contain keywords to indicate the subject matter of the image. The photo annotations provided are essential to making the photos retrievable by text-based retrieval models and allow users to formulate keyword-based queries against the photo collection.

[0006] Due to the rich nature of the image content, and the limited expressiveness of keyword-based query formulation, it is often difficult for a user to precisely formulate his information request. One reason for this difficulty is that there is often little data

associated with images. For example, an image might have data that is annotated by a user about the image. This might include the title, a description, and tags regarding the image. However, many more images lack even this rudimentary data. The data may also be noisy, meaning that the data is not relevant to the subject of the image. For example, a user might bulk upload hundreds of photos at the same time and annotate all of the photos with the same tag, without regard to content.

[0007] In addition, users often employ words in a query that may present ambiguities. Examples of word ambiguities are shown in Fig. 1. For example, a user might enter the query word “apple.” In response, the search engine might return image 101A and 101B. Image 101A shows “apple” as a fruit. Image 101B displays “apple” as a computer company logo. Unfortunately, based upon only the query word, it is not possible to determine whether the user’s intent is to find results relating to the fruit or the company logo. This type of ambiguity is referred to herein as word-sense ambiguity. In word-sense ambiguity, the dominant topic that the word may relate to is the primary sense, and a secondary topic of the word is a secondary sense. For the query “apple,” the primary sense might be “corporate logo,” and the secondary sense might be “fruit.”

[0008] In another example, a user might enter the query word “jaguar.” In response, the search engine might return image 103A and image 103B. Image 103A shows “jaguar” as an animal. Image 103B displays “jaguar” as an image of a car from the automotive manufacturer, Jaguar. The intent of the user also may not be determined based upon the query word “jaguar” without determining more detailed context.

[0009] Different ambiguities may occur as well. A determination might be made that the user intended to find images for “apple” that are related to the computer company. Under this circumstance, the images sought may refer to any one of the computer company’s products, logos, or events. In this case, the ambiguity is referred to as type-specific ambiguity.

[0010] Due to the rich nature of the image content, and the limited expressiveness of keyword-based query formulation, it is often difficult for a user to precisely formulate his information need. In the absence of disambiguating information, the user should be presented with a diverse set of images that embodies many possible interpretations of the user’s query. When presented with results reflecting multiple senses of the query, the likelihood that the user’s intention will be represented are greatly increased. While traditional information retrieval models focus on finding the most relevant document without consideration for diversity, effective image search requires results that are both

diverse and relevant. Thus, methods that provide image search results that are both relevant and diverse are highly desirable.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

[0012] FIG. 1 is a diagram displaying examples of word sense ambiguity that may be present in image search results;

[0013] FIG. 2 is a diagram displaying the steps of techniques for generating a diverse and relevant image results set, according to embodiments of the invention; and

[0014] FIG. 3 is a block diagram of a computer system on which embodiments of the invention may be implemented.

DETAILED DESCRIPTION

[0015] Techniques are described to return image search results that are both diverse and relevant, or on the contrary, highly focused and relevant. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

GENERAL OVERVIEW

[0016] Different methods may be employed to improve the diversity of image search. The goal is a method that does not harm precision for unambiguous queries while respecting the topical diversity of ambiguous queries. Methods may integrate the diversity of the search results directly into the retrieval strategy or be post-processes. Examples of post-processing are query expansion and results re-ranking. In query expansion, a user's original query is expanded to include additional words, such as synonyms or different morphological forms of a word, so that the search query matches additional documents. In results re-ranking, results are re-ranked to enrich topic coverage in retrieval results.

[0017] Different types of language modeling for retrieval may also be employed. Examples of language modeling are query likelihood and relevance models. Query

likelihood estimates the probability that a document was generated from the same distribution as the query. As used herein, a document is metadata associated with an image. Relevance models estimate a model of the query from the distribution of relevant documents in the collection. A collection is the corpus of documents from which results are obtained. Relevance models effectively add terms to the query that are related to the relevant documents. For a query that is topically unambiguous, the effect is to encourage relevant documents to be ranked higher. For topically ambiguous queries, terms related to multiple senses of the query are included in the query model, and therefore the resulting documents are more topically diverse.

[0018] Methods that return image search results that are both diverse and relevant are described herein. These methods are integrated directly into the retrieval strategy. A user initiates an image search request by submitting a query to a server that performs the search. The query comprises one or more keywords and may or may not contain ambiguous terms. The server generates a first results set by executing the original query relative to metadata associated with each image in an image corpus. In another embodiment, the original query is executed relative to only tags associated with each image in the image corpus. The server selects, from the first results set, a specified number of results that are ranked the highest. The server generates a list of terms that are derived from the metadata from each of the selected results. In an embodiment, the terms are limited to only tags from each of the selected results. The server then generates an updated query using terms from the generated list. In an embodiment, all terms in the list are included in the updated query, and each term in the updated query is weighted based upon the frequency of the occurrence of the term in the list. In another embodiment, the updated query is restricted to a specified number of terms. The terms are ranked based upon the frequency of the term in the list. The terms placed in the updated query are the specified number of highest ranked terms. Based upon the updated query, a search is performed relative to the metadata of the images to generate a second results set. The second results set is stored on a computer-readable storage medium, and then at least a subset of the second results set may be displayed to the user.

RECEIVING A QUERY AND GENERATING A FIRST RESULTS SET

[0019] When a user wishes to search for an image, he may enter one or more keywords with the intent to find images that are related to or embody concepts represented by the one or more keywords. The query is often not written in natural language, such as in the form of a question. On occasion, one or more of the keywords in

the query may present ambiguities. Examples of words with word-sense ambiguities are “apple,” “jaguar,” and “shark.” “Apple” may refer to an image of the computer company logo or a fruit. For “jaguar,” the user might be searching for an image of the automobile or an animal. For “sharks,” the user might be searching for an image of a type of fish or the logo of an NHL hockey team. It is simply not possible to determine the user’s intent without further inquiry. Generating a diverse results set solves this quandary by presenting results that correspond to various different intents.

[0020] Once the query request is received from the user, a server executes the search query request based upon metadata of the images. As used herein, “metadata” is data that describes another set of data. Metadata may refer to user-generated content, such as title, description, and tags. A title is a descriptive heading or caption for the image and is usually about three terms long. A description is a short statement that describes the image. The description is often written in natural language. Tags are keywords that may indicate or refer to any aspect of the image.

[0021] Tags associated with an image resemble a query in character, in that a tag set contains content terms not written in the sentence structure of a natural language. Tags are succinct, and users often tag photos with sets of synonyms and related terms. A user may add tags to an image upon uploading to a photo-sharing website. Other users may also add tags to particular images. For example, a user browsing a photo sharing website might particularly enjoy an image and wish to tag the image with a unique keyword so that he and others can later find the image using the keyword. Tags are not limited to synonyms or related terms, but may also refer to the genre or any aspect of the image. For example, an image of an old barn might elicit such wide ranging tags as: “barn,” “old,” “rural,” “autumn” (the season the image was taken), “nikon” (the type of the camera used), “blackandwhite” (indicating a black & white image), and “texas” (the site of the barn).

[0022] Tags may be added and rated by other users as well to provide better context. For example, an image might have a set of seven tags, but one tag might simply be the name of the camera, like “nikon.” Under this circumstance, a user browsing the image might feel that “nikon” does not reflect the content or subject matter of the image. Tags with higher ratings would be given greater importance than tags with lower ratings. Additionally, a tag that has been associated with a particular image by many different users may be given greater weight, relative to the particular image, than that given to a tag that has been associated with that image by only a few users.

[0023] Metadata may also include data that is automatically generated. For example, in the context of images, metadata might include the date the image was taken and details of the camera settings (lens, focal length, aperture). Data about the data file of the image may also be stored. Such data may include, but is not limited to, the size of the file, the type of file, and the date the file was created or modified.

[0024] The query may be executed based upon different combinations of the metadata. In an embodiment, the query is executed based upon all metadata of images (including tags). In another embodiment, the query is executed based only on tags associated with images. In other embodiments, the query is executed based on only user-generated metadata (rather than automatically generated metadata).

GENERATING A LIST OF TERMS FROM THE FIRST RESULTS SET

[0025] The server's execution of the query returns a first results set of images that are ranked based upon the relevance of the results to the query. In an embodiment, the server selects, from the first results set, a specified number of results that are ranked the highest. The specified number of results may be determined in numerous ways. In an embodiment, the specified number is a fixed number. The fixed number may be provided by an administrator and adjusted at any time. The number of results may be adjusted based upon results of previous searches. In another embodiment, the specified number may be based upon the total number of results in the first results set. For example, the specified number might vary based upon the number of results returned for a query. If up to 250 results are received, then 30 results may be selected. If 250 to 500 results are received, then 60 results may be selected. If more than 500 results are received, then 100 results may be selected. Under these circumstances, if a search yields 200 results in a result set, then the top 30 results would be selected.

[0026] In another embodiment, the specified number may also be determined based upon a percentage of the results in the first results set. The percentage may vary and may be specified by an administrator. For example, the administrator might determine that the top 10% of results from the search result set should be selected. Based upon these circumstances, if a search query yields a results set with 300 results, then the top 30 results would be selected. In yet another embodiment, the specified number is determined by the relevance score of a result. A relevance score is based on a calculation of how important the result is to the query. If a relevance score is determined to be above a particular threshold, then the result is selected. The selection process may vary based upon the implementation and may include hybrids of methods. For example, for up to 250

results in the first results set, 30 results may be selected, but for greater than 250 results in the first results set, a percentage of the results may be selected. In an embodiment, the number of results selected is less than the number of results in the first results set. In another embodiment, the number of results selected is equal to the number of results in the first results set.

[0027] In an embodiment, the server generates a list of terms that are derived from the metadata from each of the selected results (and only the selected results). Metadata for images may vary and include both user-generated content and content that is automatically generated. The server may request the metadata for each of the selected results. Once the metadata is collected, each term in the metadata is placed in a list. In an embodiment, all of the data in the metadata is used to generate the list of terms. In another embodiment, only a subset of the metadata is used to generate the list of terms. In yet another embodiment, the list is generated using only tags from each of the selected results. In another embodiment, only user-generated content in metadata is used to generate the list of terms.

[0028] The list contains many terms, and many of the terms in the list may be duplicated. To show the frequency of duplicate terms in the list, a histogram may be used. As used herein, a histogram is a graphical display of tabulated frequencies for each term in the list. The contents of the list is then used to generate an updated query.

GENERATING AN UPDATED QUERY FROM THE LIST

[0029] An updated query is then generated from the contents of the list. By including these additional terms in the updated query and executing the updated query, the server derives a more diverse result set than the first results set, which was based upon only the original query. In an embodiment, all terms that comprise the list are used in the updated query. When all terms are included, the terms within the updated query are weighted. The terms may be weighted based upon the frequency of a particular term in the list. For example, the list might have a total of ten terms with four terms in the stated frequencies for the original query “beetle”: “insect” (5 times) “car” (3 times) “band” (1 time) “comic” (1 time). Each term is weighted by dividing the frequency of the term in the list by the total number of terms in the list. Thus, for the term “insect,” the weighting is calculated as $5/10 = 0.5$. Under these circumstances, the updated query would appear: insect (0.5), car (0.3), band (0.1), comic (0.1). Terms that appear more frequently in the list are given a greater weighting, and more importance and influence, in the query. In the case where

only tags are used to generate the list, then ratings of the tags and/or the number of users that submitted those tags may also be taken into account when calculating weightings.

[0030] In an embodiment, the updated query may be generated by taking a specified number of the most frequently occurring terms from the list. In an embodiment, the number may be specified by an administrator and is a set number regardless of the number of unique terms in the list. In another embodiment, the specified number is determined by a percentage of the number of unique terms in the list.

[0031] In an embodiment, the updated query may be generated by taking terms from the list that occur at a frequency ratio that is above a threshold ratio. A frequency ratio is determined for each term in the list. Frequency ratio of a particular term is calculated by dividing the number of occurrences of the particular term by the total number of terms in the list. For example, if the term “apple” appeared 50 times in the list, and there were a total of 300 terms (where each duplicate term is counted separately) in the list, then the frequency ratio of “apple” would be $50/300 = 0.167$. An administrator might specify that a threshold of a frequency ratio of 0.15 must be satisfied for a term to be included in the updated query. Thus, each term with a frequency ratio equal to or greater than 0.15 would be included in the updated query.

[0032] When the updated query is generated, the original query terms may or may not be present in the updated query. In an embodiment, the original query keywords are included in the updated query. For example, in such an embodiment, if a user enters the original search query “apple”, then “apple” would also always appear in the updated query generated. The original query keyword may be assigned a specified weighting to be used in the updated query. This weighting may be greater than the weighting of any other term in the updated query, or may be a weighting lower than the weighting of any other term in the updated query. In another embodiment, all original query keywords are excluded from the updated query. In yet another embodiment, original query terms only appear in the updated query if the original query term appears sufficiently frequent in the list.

GENERATING THE SECOND RESULTS SET FROM THE UPDATED QUERY

[0033] The server generates a second results set by executing the updated query. The results set produced by executing the updated query is ranked by relevance. In an embodiment, the updated query is executed relative to all of the metadata of the images in the image corpus (which may include images not in the first results set). In another embodiment, the updated query is executed relative to only the tags of images in the

image corpus. In yet another embodiment, the updated query is executed relative only to metadata that is user-generated for the images in the image corpus.

[0034] The second results set may be stored on volatile or non-volatile computer readable storage medium. At least a portion of the second results set (comprising a relevance-ranked set of images with related metadata) may then be displayed to the user as the results of the original query, since the first results set is never shown to the user.

ACHIEVING DESIRED IMAGE SEARCH RESULTS

[0035] Using the retrieval methods detailed above, the unit of retrieval to which the queries are executed may be varied to achieve a desired result in image search. The desired result may be increased diversity of the image search results set or more precision of the image result set. Particular results may provide a more diverse image search result based upon the unit of retrieval, and other particular results may provide a more focused image search result if a different unit of retrieval is used. The unit of retrieval may also vary (dual index approach) where the original query uses a particular unit of retrieval and the updated query uses a different particular unit of retrieval. In an embodiment, results are generated where the unit of retrieval is restricted to the tags of the metadata of the images in the image corpus. In an embodiment, this may be used with any type of language model employed. For example, the original query and updated query might be executed relative to only the tags of the metadata of images. The contents of the list may also be restricted to only tags of the metadata.

[0036] In an embodiment, a dual index system is employed. In a dual index system, the original query is executed relative to only tags in the metadata of the images in the image corpus and the updated query is executed relative to all metadata of the images in the image corpus. In another embodiment of the dual index system, the original query is executed relative to all metadata of the images in the image corpus and the updated query is executed relative to only tags of the metadata of the images in the image corpus. The contents of the list may be restricted to only tags of the metadata or encompass all metadata of the images.

ILLUSTRATED STEPS

[0037] Fig. 2 is a diagram displaying steps of techniques for generating a diverse and relevant image results set. In step 201, a user submits an original query that is received by a server to search for images. In step 203, the server generates a first result set by executing the original query. The query is executed relative to the metadata of images in

an image corpus. In other embodiments, the query may be executed relative to only tags of those images. Once the first result set is received, the server selects a specified number of the highest ranked results. In step 205, the server generates a list that comprises the terms from the metadata of each of the selected results. Once the list is generated, then one of two steps may occur, as illustrated in alternative steps 207A and 207B. In step 207A, the server generates an updated query that includes one instance of each term that is in the list. Under step 207A, each term in the updated query is weighted based upon the frequency of the particular term in the list. Alternatively, in step 207B, an updated query is generated that only includes a specified number of terms that appear most frequently in the list. In other embodiments, the terms in the updated query of 207B may also be weighted. Regardless of whether step 207A or 207B is performed, control then passes to step 209. In step 209, the server generates a second results set by executing the updated query relative to the metadata of the images in the image corpus (which may include images that were not in the first results set). In other embodiments, the query may be executed relative to only the tags of images. The second results set is stored and at least a portion of the second results set is displayed to the user as the results of the original query. Consequently, a more diverse set of images, representing multiple different possible meanings of ambiguous query terms in the original query, is provided to the user.

HARDWARE OVERVIEW

[0038] Figure 3 is a block diagram that illustrates a computer system 300 upon which an embodiment of the invention may be implemented. Computer system 300 includes a bus 302 or other communication mechanism for communicating information, and a processor 304 coupled with bus 302 for processing information. Computer system 300 also includes a main memory 306, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 302 for storing information and instructions to be executed by processor 304. Main memory 306 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 304. Computer system 300 further includes a read only memory (ROM) 308 or other static storage device coupled to bus 302 for storing static information and instructions for processor 304. A storage device 310, such as a magnetic disk or optical disk, is provided and coupled to bus 302 for storing information and instructions.

[0039] Computer system 300 may be coupled via bus 302 to a display 312, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device

314, including alphanumeric and other keys, is coupled to bus 302 for communicating information and command selections to processor 304. Another type of user input device is cursor control 316, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 304 and for controlling cursor movement on display 312. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

[0040] The invention is related to the use of computer system 300 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 300 in response to processor 304 executing one or more sequences of one or more instructions contained in main memory 306. Such instructions may be read into main memory 306 from another machine-readable medium, such as storage device 310. Execution of the sequences of instructions contained in main memory 306 causes processor 304 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

[0041] The term “machine-readable medium” as used herein refers to any medium that participates in providing data that causes a machine to operation in a specific fashion. In an embodiment implemented using computer system 300, various machine-readable media are involved, for example, in providing instructions to processor 304 for execution. Such a medium may take many forms, including but not limited to storage media and transmission media. Storage media includes both non-volatile media and volatile media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 310. Volatile media includes dynamic memory, such as main memory 306. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 302. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications. All such media must be tangible to enable the instructions carried by the media to be detected by a physical mechanism that reads the instructions into a machine.

[0042] Common forms of machine-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory

chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

[0043] Various forms of machine-readable media may be involved in carrying one or more sequences of one or more instructions to processor 304 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 300 can receive the data on the telephone line and use an infra-red transmitter to convert the data to an infra-red signal. An infra-red detector can receive the data carried in the infra-red signal and appropriate circuitry can place the data on bus 302. Bus 302 carries the data to main memory 306, from which processor 304 retrieves and executes the instructions. The instructions received by main memory 306 may optionally be stored on storage device 310 either before or after execution by processor 304.

[0044] Computer system 300 also includes a communication interface 318 coupled to bus 302. Communication interface 318 provides a two-way data communication coupling to a network link 320 that is connected to a local network 322. For example, communication interface 318 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 318 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 318 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0045] Network link 320 typically provides data communication through one or more networks to other data devices. For example, network link 320 may provide a connection through local network 322 to a host computer 324 or to data equipment operated by an Internet Service Provider (ISP) 326. ISP 326 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 328. Local network 322 and Internet 328 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 320 and through communication interface 318, which carry the digital data to and from computer system 300, are exemplary forms of carrier waves transporting the information.

[0046] Computer system 300 can send messages and receive data, including program code, through the network(s), network link 320 and communication interface 318. In the

Internet example, a server 330 might transmit a requested code for an application program through Internet 328, ISP 326, local network 322 and communication interface 318.

[0047] The received code may be executed by processor 304 as it is received, and/or stored in storage device 310, or other non-volatile storage for later execution. In this manner, computer system 300 may obtain application code in the form of a carrier wave. In the foregoing specification, embodiments of the invention have been described with reference to numerous specific details that may vary from implementation to implementation. Thus, the sole and exclusive indicator of what is the invention, and is intended by the applicants to be the invention, is the set of claims that issue from this application, in the specific form in which such claims issue, including any subsequent correction. Any definitions expressly set forth herein for terms contained in such claims shall govern the meaning of such terms as used in the claims. Hence, no limitation, element, property, feature, advantage or attribute that is not expressly recited in a claim should limit the scope of such claim in any way. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

EXAMPLES

[0048] In an embodiment, a method to generate a results set of images comprises or a computer-readable storage medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform: receiving an original query from a user to search a plurality of images, each of which is associated with metadata; generating a first results set based on executing the original query relative to tags from the metadata associated with each image in said plurality of images; selecting, from the first results set, a specified number of results ranked highest in the first results set; generating a list comprising terms from tags from the metadata associated with each result of the specified number of results; for each term in the list, calculating a weight for that term based on a frequency of occurrence of that term in the list; generating an updated query comprising terms that are in the list and that are weighted based on said weights; generating a second results set based on executing the updated query relative to all metadata associated with each image in said plurality of images; and storing the second results set on a non-volatile or volatile computer readable storage medium.

[0049] In an embodiment, a method to generate a results set of images comprises or a computer-readable storage medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform:

receiving an original query from a user to search a plurality of images, each of which is associated with metadata; generating a first results set based on executing the original query relative to tags from the metadata associated with each image in said plurality of images; selecting, from the first results set, a specified number of results ranked highest in the first results set; generating a list comprising terms from tags from the metadata associated with each result of the specified number of results; for each term in the list, calculating a weight for that term based on a frequency of occurrence of that term in the list; generating an updated query comprising terms that are in the list and that are weighted based on said weights; generating a second results set based on executing the updated query relative to tags from the metadata associated with each image in said plurality of images; and storing the second results set on a non-volatile or volatile computer readable storage medium.

[0050] In an embodiment, a method to generate a results set of images comprises or a computer-readable storage medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform: receiving an original query from a user to search a plurality of images, each of which is associated with metadata; generating a first results set based on executing the original query relative to metadata associated with each image in said plurality of images; selecting, from the first results set, a specified number of results ranked highest in the first results set; generating a list comprising terms from metadata associated with each result of the specified number of results; for each term in the list, calculating a weight for that term based on a frequency of occurrence of that term in the list; generating an updated query comprising terms that are in the list and that are weighted based on said weights; generating a second results set based on executing the updated query relative to metadata associated with each image in said plurality of images; and storing the second results set on a non-volatile or volatile computer readable storage medium.

[0051] In an embodiment, a method or a computer-readable storage medium wherein the list comprises terms from only tags, including tags that were associated with particular images by users other than users that uploaded the particular images.

[0052] In an embodiment, a method or a computer-readable storage medium wherein the specified number of results is based upon the number of relevant results generated by the original query.

[0053] In an embodiment, a method or a computer-readable storage medium wherein the specified number of results is based upon a percentage of a quantity of relevant results generated by the original query.

[0054] In an embodiment, a method or a computer-readable storage medium wherein metadata comprises title, description, and tags.

[0055] In an embodiment, a method or a computer-readable storage medium wherein the updated query comprises terms contained in the original query.

[0056] In an embodiment, a method or a computer-readable storage medium wherein the updated query does not comprise terms contained in the original query.

[0057] In an embodiment, a method or a computer-readable storage medium wherein the specified number of results is less than all of the results in the first results set.

[0058] In an embodiment, a method to generate a results set of images comprises or a computer-readable storage medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform: receiving an original query from a user to search a plurality of images, each of which is associated with metadata; generating a first results set based on executing the original query relative to metadata associated with each image in said plurality of images; selecting, from the first results set, a first specified number of results ranked highest in the first results set; generating a list comprising terms from metadata associated with each result of the specified number of results; selecting, from the list, a second specified number of terms that occur most frequently in the list; generating an updated query comprising the second specified number of terms selected from the list; generating a second results set based on executing the updated query relative to metadata associated with each image in said plurality of images; and storing the second results set on a non-volatile or volatile computer readable storage medium.

[0059] In an embodiment, a method or a computer-readable storage medium wherein the list comprises terms from only tags, including tags that were associated with particular images by users other than users that uploaded the particular images.

[0060] In an embodiment, a method or a computer-readable storage medium wherein the specified number of results is based upon the number of relevant results generated by the original query.

[0061] In an embodiment, a method or a computer-readable storage medium wherein the specified number of results is based upon a percentage of a quantity of relevant results generated by the original query.

[0062] In an embodiment, a method or a computer-readable storage medium wherein metadata comprises title, description, and tags.

CLAIMS

What is claimed is:

1. A computer-implemented method to generate a results set of images, comprising:
receiving an original query from a user to search a plurality of images, each of
which is associated with metadata;
generating a first results set based on executing the original query relative to tags
from the metadata associated with each image in said plurality of images;
selecting, from the first results set, a specified number of results ranked highest in
the first results set;
generating a list comprising terms from tags from the metadata associated with
each result of the specified number of results;
for each term in the list, calculating a weight for that term based on a frequency of
occurrence of that term in the list;
generating an updated query comprising terms that are in the list and that are
weighted based on said weights;
generating a second results set based on executing the updated query relative to all
metadata associated with each image in said plurality of images; and
storing the second results set on a non-volatile or volatile computer readable
storage medium.
2. A computer-implemented method to generate a results set of images, comprising:
receiving an original query from a user to search a plurality of images, each of
which is associated with metadata;
generating a first results set based on executing the original query relative to tags
from the metadata associated with each image in said plurality of images;
selecting, from the first results set, a specified number of results ranked highest in
the first results set;
generating a list comprising terms from tags from the metadata associated with
each result of the specified number of results;
for each term in the list, calculating a weight for that term based on a frequency of
occurrence of that term in the list;
generating an updated query comprising terms that are in the list and that are
weighted based on said weights;

- generating a second results set based on executing the updated query relative to tags from the metadata associated with each image in said plurality of images; and
- storing the second results set on a non-volatile or volatile computer readable storage medium.
3. A computer-implemented method to generate a results set of images, comprising:
- receiving an original query from a user to search a plurality of images, each of which is associated with metadata;
- generating a first results set based on executing the original query relative to metadata associated with each image in said plurality of images;
- selecting, from the first results set, a specified number of results ranked highest in the first results set;
- generating a list comprising terms from metadata associated with each result of the specified number of results;
- for each term in the list, calculating a weight for that term based on a frequency of occurrence of that term in the list;
- generating an updated query comprising terms that are in the list and that are weighted based on said weights;
- generating a second results set based on executing the updated query relative to metadata associated with each image in said plurality of images; and
- storing the second results set on a non-volatile or volatile computer readable storage medium.
4. The method of Claim 3, wherein the list comprises terms from only tags, including tags that were associated with particular images by users other than users that uploaded the particular images.
5. The method of Claim 3, wherein the specified number of results is based upon the number of relevant results generated by the original query.
6. The method of Claim 3, wherein the specified number of results is based upon a percentage of a quantity of relevant results generated by the original query.
7. The method of Claim 3, wherein metadata comprises title, description, and tags.

8. The method of Claim 3, wherein the updated query comprises terms contained in the original query.
9. The method of Claim 3, wherein the updated query does not comprise terms contained in the original query.
10. The method of Claim 3, wherein the specified number of results is less than all of the results in the first results set.
11. A computer-implemented method to generate a results set of images, comprising:
 - receiving an original query from a user to search a plurality of images, each of which is associated with metadata;
 - generating a first results set based on executing the original query relative to metadata associated with each image in said plurality of images;
 - selecting, from the first results set, a first specified number of results ranked highest in the first results set;
 - generating a list comprising terms from metadata associated with each result of the specified number of results;
 - selecting, from the list, a second specified number of terms that occur most frequently in the list;
 - generating an updated query comprising the second specified number of terms selected from the list;
 - generating a second results set based on executing the updated query relative to metadata associated with each image in said plurality of images; and
 - storing the second results set on a non-volatile or volatile computer readable storage medium.
12. The method of Claim 11, wherein the list comprises terms from only tags, including tags that were associated with particular images by users other than users than uploaded the particular images.
13. The method of Claim 11, wherein the specified number of results is based upon the number of relevant results generated by the original query.


14. The method of Claim 11, wherein the specified number of results is based upon a percentage of a quantity of relevant results generated by the original query.
15. The method of Claim 11, wherein metadata comprises title, description, and tags.
16. A computer-readable storage medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to:
 - receive an original query from a user to search a plurality of images, each of which is associated with metadata;
 - generate a first results set based on executing the original query relative to metadata associated with each image in said plurality of images;
 - select, from the first results set, a specified number of results ranked highest in the first results set;
 - generate a list comprising terms from metadata associated with each result of the specified number of results;
 - for each term in the list, calculate a weight for that term based on a frequency of occurrence of that term in the list;
 - generate an updated query comprising terms that are in the list and that are weighted based on said weights;
 - generate a second results set based on executing the updated query relative to metadata associated with each image in said plurality of images; and
 - store the second results set on a non-volatile or volatile computer readable storage medium.
17. The computer-readable storage medium of Claim 16, wherein the list comprises terms from only tags, including tags that were associated with particular images by users other than users than uploaded the particular images.
18. The computer-readable storage medium of Claim 16, wherein the specified number of results is based upon the number of relevant results generated by the original query.

19. The computer-readable storage medium of Claim 16, wherein the specified number of results is based upon a percentage of a quantity of relevant results generated by the original query.
20. The computer-readable storage medium of Claim 16, wherein metadata comprises title, description, and tags.
21. The computer-readable storage medium of Claim 16, wherein the updated query comprises terms contained in the original query.
22. The computer-readable storage medium of Claim 16, wherein the updated query does not comprise terms contained in the original query.
23. The computer-readable storage medium of Claim 16, wherein the specified number of results is less than all of the results in the first results set.
24. A computer-readable storage medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to:
 - receive an original query from a user to search a plurality of images, each of which is associated with metadata;
 - generate a first results set based on executing the original query relative to metadata associated with each image in said plurality of images;
 - select, from the first results set, a first specified number of results ranked highest in the first results set;
 - generate a list comprising terms from metadata associated with each result of the specified number of results;
 - select, from the list, a second specified number of terms that occur most frequently in the list;
 - generate an updated query comprising the second specified number of terms selected from the list;
 - generate a second results set based on executing the updated query relative to metadata associated with each image in said plurality of images; and
 - store the second results set on a non-volatile or volatile computer readable storage medium.


25. The computer-readable storage medium of Claim 24, wherein the list comprises terms from only tags, including tags that were associated with particular images by users other than users than uploaded the particular images.
26. The method of Claim 24, wherein the specified number of results is based upon the number of relevant results generated by the original query.
27. The method of Claim 24, wherein the specified number of results is based upon a percentage of a quantity of relevant results generated by the original query.
28. The method of Claim 24, wherein metadata comprises title, description, and tags.

FIG. 1
(prior art)


101A



101B



103A



103B





FIG. 2

FIG. 3

