

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
18 May 2006 (18.05.2006)

PCT

(10) International Publication Number
WO 2006/053264 A1(51) International Patent Classification:
G06F 17/30 (2006.01)

(21) International Application Number:

PCT/US2005/041074

(22) International Filing Date:

10 November 2005 (10.11.2005)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

60/627,254	11 November 2004 (11.11.2004)	US
60/627,189	11 November 2004 (11.11.2004)	US
11/150,369	10 June 2005 (10.06.2005)	US
11/150,045	10 June 2005 (10.06.2005)	US
11/169,521	28 June 2005 (28.06.2005)	US

(71) Applicant (for all designated States except US): **YAHOO! INC.** [US/US]; D-274, 701 First Avenue, Sunnyvale, California 94089 (US).

(72) Inventors; and


(75) Inventors/Applicants (for US only): **CARSON, Chad**[US/US]; 990 Governors Bay Drive, Rewood City, California 94065 (US). **COOK, Douglas M.** [US/US]; 212 Dorland Street, San Francisco, California 94144 (US). **RAVINARAYANAN, Kalpana** [IN/US]; 1057 Marigold Court, Sunnyvale, California 94086 (US).(74) Agents: **HICKMAN, Brian, D.** et al.; Hickman Palermo Truong & Becker LLP, 2055 Gateway Place, Suite 550, San Jose, CA 95110-1089 (US).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH,

[Continued on next page]

(54) Title: ACTIVE ABSTRACTS



(57) Abstract: Techniques are provided for generating search results that include, instead of or in addition to the traditional abstract links, other links to content that the user of a search engine may consider helpful. These other links are referred to herein as "non-traditional abstract links". In general, the search engine generates non-traditional abstract links that attempt to anticipate the "next move" by the search engine user. By selecting a non-traditional abstract link from the abstract of a matching resource, the user is able to navigate from the search result listing directly to locations other than the top of a matching resource.



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, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

ACTIVE ABSTRACTS

FIELD OF THE INVENTION

[0001] The present invention relates to generating search results and, more specifically, generating abstracts that contain non-traditional abstract links.

BACKGROUND

[0002] A search engine is a computer program that helps a user to locate information. To locate information on a particular topic, a user can submit to a search engine one or more search query terms related to the topic. In response, the search engine executes the search query and generates information about the results of the search. The information about the results of the search, referred to herein as the "search results", usually contains a list of the resources that satisfy the search query. The resources identified in the search results are referred to herein as "matching resources".

[0003] While search engines may be applied in a variety of contexts, search engines are especially useful for locating resources that are accessible through the Internet. Resources may include files whose content is composed in a page description language such as Hypertext Markup Language (HTML). Such files are typically called pages. Using a web browser, pages may be retrieved by selecting HTML links that contain the Universal Resource Locators (URLs) of the pages.

[0004] Depending on the query terms used and the number of pages that contain those query terms, search results may contain so many matching resources that a user may be overwhelmed when trying to determine which matching resources to investigate further. To assist a user in selecting one or more matching resources from a list, the search results may include a short description, or abstract, for each matching resource. By reading the abstract for a given matching resource, a user should be able to better determine whether the matching resource merits further investigation. Abstracts should be relatively short, so that a user may quickly judge the relevance of matching resources listed in the search results.

[0005] Various efforts have been made to improve the content of the abstracts that are presented to users. For example, one technique to improve the content of abstracts is described in U.S. Patent Application No. 10/365,273, entitled "GENERATING DESCRIPTIONS OF MATCHING RESOURCES BASED ON THE KIND, QUALITY, AND RELEVANCE OF AVAILABLE SOURCES OF INFORMATION ABOUT THE MATCHING RESOURCES", the contents of which are incorporated herein by reference. That application describes a technique for selecting the process which generates the abstract for a web page, from a plurality of available processes,

based on the kind, quality, and relevance of the available sources of information about the page.

[0010] Another technique to improve the content of the abstracts is described in U.S. Patent No. 6,711,567, entitled "DELIVERING NON-DEFAULT ITEMS IN ASSOCIATION WITH SEARCH RESULTS", the contents of which are incorporated by this reference. The "NON-DEFAULT" patent describes a technique for delivering search results pages to the users of a search engine, where one or more search result listings on the search results pages include items that are not included by default. The non-default items may take many forms, including but not limited to images, banners, controls, animations, and even applets in the JAVA language. The non-default items may be designed to entice the search engine users to select the search result listing with which they are associated over possibly higher-ranked search result listings that contain only default items.

[0011] The party that controls a web site may be pleased if the use of non-default items in the search results entices more users to visit the web site. However, the search engine user's experience is not necessarily improved. Therefore, it is desirable to provide techniques for making the contents of search result listing more useful to a user, thereby facilitating the user's search for useful information.

[0012] 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.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013] 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:

[0014] FIG. 1 is a flowchart that illustrates a two-phase process for selecting which non-traditional abstract links to include in an abstract, according to an embodiment of the invention;

[0015] FIGs. 2-9 are diagrams that illustrate examples of abstracts that may be generated using the techniques described herein; and

[0016] FIG. 10 is a block diagram of a computer system that may be used to implement the techniques described herein.

DETAILED DESCRIPTION

[0017] 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.

NON-TRADITIONAL ABSTRACT LINKS

[0018] Typically, search results include, for each matching resource, an abstract and a link. When selected, the link for a matching resource causes the browser to request the matching resource. When the matching resource is a web page, selection of the link causes the web page to be retrieved into the browser. If the web page is longer than will fit in the browser's window, then the top of the web page is displayed, and controls are provided that allow a user to scroll down to see more of the web page. Such links are referred to herein as "traditional abstract links".

[0019] Referring to FIG. 2, it illustrates a search result listing that includes two abstracts 200 and 202. The abstract 200 includes a heading 204, which, when selected, activates a traditional abstract link to retrieve the matching resource. Abstract 200 also displays the URL 206 of the matching resource. Similar to heading 204, when URL 206 is selected, the matching resource is retrieved. Thus, both heading 204 and URL 206 are associated with the traditional abstract link of the matching resource associated with abstract 200.

[0020] Techniques shall be described hereafter for generating search results that include, instead of or in addition to the traditional abstract links, other links to contents that the user may consider helpful. These other links are referred to herein as "non-traditional abstract links". In general, the search engine generates non-traditional links in an attempt to anticipate the "next move" by the search engine user. By selecting a non-traditional abstract link from the abstract of a matching resource, the user is able to navigate from the search result listing directly to locations other than the top of a matching resource. Various examples of non-traditional abstract links, and techniques for generating abstracts that contain such links, shall be described in detail hereafter.

[0021] Referring again to FIG. 2, abstract 200 includes an address 208 and the word "Photos" 210. Each of these items is "active" in that selecting the item activates a link. However, unlike heading 204 and URL 206, selection of address 208 or "Photos" 210 will activate a non-traditional abstract link, as shall be described in greater detail below.

[0022] When displayed as part of the abstract of a matching resource, non-traditional abstract links may or may not be visibly separate from the rest of the abstract. For example, some non-traditional abstract links (such as the link associated with address 212 in abstract 202) may simply appear as text, within the textual body of the abstract. On the other hand, some non-traditional abstract links (such as those associated with address 208 and “Photos” 210) may appear separate from, but in the vicinity of, the textual body of the abstract. Yet other non-traditional abstract links may be displayed in a manner that does not visibly convey the relationship between the non-traditional abstract links and the textual body of the abstract.

ACTIVATING NON-TRADITIONAL ABSTRACT LINKS

[0023] When a link is included in a web page, the link itself may or may not be visible to a viewer of the web page. For example, a web page display may anchor text which, when selected, causes the corresponding link to be activated. The text that causes the link to be selected may be entirely different than the text of the link itself. For example, the link may be www.findanewcar.com but the text that activates the link may be “My Favorite Car Site”. Similarly, the web page may contain a non-textual control, such as a button or image which, when selected, activates the link. Thus, in the descriptions that follow, an abstract that includes non-traditional search links does not necessarily display the text of the link itself.

[0024] As mentioned above, non-traditional abstract links may be included in the abstract associated with a matching resource to provide the search engine user direct access to potentially helpful information. To view the information associated with a non-traditional search result link, the user must perform some action to activate the non-traditional abstract link.

[0025] According to one embodiment, a non-traditional abstract link is activated in the same manner as a traditional abstract link. Specifically, the user selects the link by clicking anchor text that is displayed in association with the link. Most browsers will display such anchor text in a manner that distinguishes the anchor text from other text that is not associated with a link. For example, many browsers will display text that is associated with a link using a blue underlined font.

[0026] In alternative embodiments, a non-traditional abstract link is activated in a non-traditional manner. For example, in one embodiment, a non-traditional abstract link is activated by merely “hovering” the cursor over the text associated with the link. Specifically, if the location of the cursor remains in a region of the display screen associated with the link for more than a predetermined threshold of time, the link is activated.

[0027] As another example of how a non-traditional abstract link may be activated, hovering over the link may cause a user interface control to appear next to the text. Once the user interface control appears, the user may “hover” the cursor over the user interface control to activate the link. Alternatively, after the user interface control appears, the user may have to click on the user interface control to activate the link.

MID-DOCUMENT LINKS

[0028] As mentioned above, traditional abstract links allow the user to navigate from the search results directly to the top of a matching resource, such as a web page. Unfortunately, the part of the web page in which the user is interested may be far from the top of the web page. Consequently, even though the traditional abstract link helps the user to retrieve a web page that contains relevant content, finding the relevant content within the web page may still be cumbersome.

[0029] According to one embodiment, search results are generated with mid-document links. Selection of a mid-document link causes a browser not only to retrieve a matching resource, but to jump to a particular location within the resource. Including mid-document links in the abstract for a matching resource is particularly useful if the matching links will take the user to the location of content that will be relevant to the user.

[0030] The terms used in the search query performed by a user is an indication of what the user considers to be relevant. Thus, the portion of a matching document that contains those search terms has a relatively high likelihood of being relevant to the user. Therefore, according to one embodiment, the abstract associated with a matching resource is generated to include a mid-document link to the portion of the matching resource that contains the terms that were used in the search query.

[0031] A mid-document link identifies an “anchor” within a web page. When a browser retrieves a web page based on a mid-document link, the browser locates the portion of the web page to display by searching for the anchor identified in the mid-document link. However, a web page may not have an anchor embedded at the location contains the terms that match a search query. Therefore, in one embodiment, the search engine:

- identifies a matching resource based on the search query,
- embeds an anchor in the matching resource at the location within the matching resource that contains terms that match the search query,

- generates a mid-document link which, when selected, will cause a browser to retrieve the matching resource and to automatically jump to the location of the embedded anchor, and
- sends, as part of the search results, an abstract for the matching resource, where the abstract includes the mid-document link.

[0032] U.S. Patent Application No. 11/150,369, entitled DOCUMENT NAVIGATION BY ACTIVE ABSTRACTS, filed on June 10, 2005 (attorney docket no. 50269-0642), which has been incorporated herein by reference, describes in greater detail the generation and use of mid-document links within abstracts that are provided with search result listings.

[0033] In the example given above, the mid-document links are links to the location, within the matching resource, of terms that matched the search query. However, mid-document links are not limited to links to the location of matching search terms. Rather, mid-document links may be included in the abstract of a matching resource to allow the search engine user to immediately jump to the location within the matching resource of content of interest. Thus, in one embodiment, the abstract of a web page that has section headings may be constructed to include a mid-document link to each section heading.

[0034] Referring again to FIG. 2, address 208 and "Photos" 210 may be mid-document links. Abstract 200 may be associated with a matching resource that includes an address and photos. In this case, address 208 may have been extracted by the search engine from the matching resource. In constructing the abstract 208, the search engine uses the address 208 as anchor text for the mid-document link to the location, within the matching resource, from which address 208 was extracted. Similarly, Photos 210 may be the anchor text for a mid-document link to the location, within the matching resource, of the digital images.

[0035] As described in more detail in the related applications, a cached version of a matching resource may be used to implement the mid-document links. The cached version of the matching resource may be modified by the search engine to embed anchors for the mid-document links without modifying the version of the matching resource that resides at its original source.

PROMOTED LINKS

[0036] According to one embodiment, techniques are provided for including "promoted links" within the abstracts that are provided with search result listings. Unlike a traditional search result link, selection of a promoted link does not cause retrieval of the matching resource associated with the abstract in which the promoted link resides. Rather, a promoted

link is a link that either (1) is contained within the matching resource, or (2) is contained within a resource to which the matching resource is directly or indirectly linked.

[0037] For example, assume that a web page A includes a link to a web page B, and that web page B includes a link to a web page C. Assume further that web page A is a matching resource for a particular search query. The abstract for web page A in the search results for that particular query may include a link to web page B and/or a link to web page C. In this example, the links to web pages B and C, within the abstract for web page A, are promoted links.

[0038] U.S. Patent Application No. 11/150,045, entitled SEARCH SYSTEM PRESENTING ACTIVE ABSTRACTS INCLUDING LINKED TERMS, filed on June 10, 2005 (attorney docket no. 50269-0641), which has been incorporated herein by reference, describes in greater detail the generation and use of promoted links within abstracts that are provided with search result listings.

[0039] The search engine may incorporate sophisticated logic to determine which links to promote to an abstract. For example, assume that a particular page about digital cameras is a matching resource in a search for digital cameras. In this case, the matching resource may itself have several links to web pages about specific digital cameras. According to one embodiment, the search engine determines which of those digital cameras are the "most popular", and promotes only the links to the most popular of the digital cameras.

[0040] As another example, assume that a web page about specific digital camera is a matching resource in a search for digital cameras. In this case, the matching resource may contain a link to a web page that has several links to web pages about specific digital cameras. As in the example give above, the search engine may determine which of those digital cameras are the "most popular", and promote only the links to the most popular of the digital cameras. Unlike the previous example, in this example, the promoted links are to web pages that are two levels below the resource that matched the actual search.

[0041] Another kind of "promoted link" is a link to a more general page on the same site. For example, the abstract for a page, from a particular site, about a specific digital camera could include a link to the page, from that particular site, about digital cameras in general, or a page about other digital cameras within the same price range or from the same manufacturer.

[0042] Referring again to FIG. 2, address 208 may not appear within the matching resource at all. Companies frequently put their address on a "Contact Us" page that is one or more levels below their home page. Thus, the non-traditional abstract link associated with

address 208 may be included in abstract 208 as a result of being promoted from the Contact Us page.

SIMILAR SUBJECT MATTER LINKS

[0043] According to one embodiment, techniques are provided for including “similar subject matter links” within the abstracts that are provided with search result listings. Unlike a promoted link, a similar subject matter link need not be a link within a matching resource, nor within any resource to which the matching resource is linked. For example, the search result listing for a particular search query may include an abstract for a particular web page about Australia. That particular web page may not include any links at all. However, due to the similarity of subject matter, the search engine may include, in the abstract for the web page about Australia, a link to another web page about kangaroos.

[0044] As another example, the search result listing for a particular search query may include an abstract for a web page that describes “bar tables”. Frequently, people interested in information about bar tables will also be interested in information about “bar stools”. Thus, the search engine may include, within the abstract of the web page for bar tables, a non-traditional abstract link to a different page for bar stools, even though the page for bar tables does not include a link to that particular page about bar stools.

[0045] Referring again to FIG. 2, the Nite Inn web page associated with abstract 200 may not actually have any photographs. Thus, “Photos” 210 may be the anchor text for a similar subject matter link which, when activated, retrieves a web page, other than the Nite Inn web page associated with abstract 200, that includes photographs of the Nite Inn.

VALUE ADDED SERVICE LINKS

[0046] According to one embodiment, techniques are provided for including “value added service links” within the abstracts that are provided with search result listings. Similar to a similar subject matter link, a value added service link need not be a link within a matching resource. Rather, value added service links are links which, when selected, invoke services that generate output based, in part, on information relating to the matching resource.

[0047] For example, the search result listing for a particular search query may include an abstract for a particular web page about a hotel. The hotel’s web page may include the address of the hotel. Under these circumstances, the search engine may include, in the abstract for the hotel’s web page, a link which, when activated, sends the hotel’s address to a web site that uses the hotel’s address to generate a map of the vicinity around the hotel’s address.

[0048] As another example, the hotel's web page may include a phrase that is in a foreign language, and that phrase may be included in the abstract for the hotel's web page. Under these circumstances, the search engine may include, in the abstract for the hotel's web page, a link which, when activated, sends the foreign language phrase to a web site that translates the foreign language phrase to English. In this example, the search engine may construct the abstract in such a way that selection of the foreign language phrase activates the link to the translation service.

[0049] In addition to or instead of value added service links, the search engine may include, in the abstract, value added information. Value added information is similar to value added links, except that it is not presented as a link which, when activated, retrieves additional information. Value added information may include, for example, phone numbers and addresses, and controls associated with different ways to sort/filter the results page. The value added information may or may not exist in the content of the matching resource. For example, a company's home page may not have a phone number for the company. However, the phone number of a company may be useful to a user. In one embodiment, the search engine provides information about the company to a phone directory service, and receives from the directory service a phone number for the company. When constructing the abstract for the company, the search engine may include the phone number in the abstract.

[0050] The search engine may invoke the phone directory service to obtain the phone number in response to receiving the search query. Alternatively, the search engine may invoke the phone directory service ahead of time, and store the phone number in metadata associated with the company web page. In the latter case, the search engine may read the phone number from the metadata when constructing the abstract in response to a particular query.

[0051] Referring again to FIG. 2, address 208 may be associated with a value added service link which, when selected, sends address information to a service that will generate a map of the vicinity around the address.

RELATED SEARCH LINKS

[0052] According to one embodiment, the search engine may include, in the abstract for a particular web page, a link which, when activated, initiates another search. The search that is initiated by such an embedded "related search link" may use terms extracted from the web page represented by the abstract, and/or terms not necessarily from the web page, but which are related to the content of the web page. For example, the search engine may include, in the

abstract for a hotel's web page, a link which, when activated, initiates another search for reviews of that particular hotel.

[0053] The search terms used by the related search link may take into account both the original search terms (which resulted in the display of the abstract that contains the search link) and terms related to the content of the web page represented by the abstract. For example, assume that the original search used the terms "hotels San Jose California". The matching resources for that search may include a web page for a Hotel X. In constructing the abstract for the Hotel X web page, the search engine may include a link which, when activated, initiates a search based on the terms "Hotel X San Jose California reviews".

RESPONDING TO SELECTION OF NON-TRADITIONAL ABSTRACT LINKS

[0054] As mentioned above, instead of or in addition to the traditional abstract link, the search engine may embed non-traditional abstract links within the abstract of a matching resource so that the user of the search engine may directly retrieve relevant information. When such a link is selected, the corresponding information may be presented to the user in a variety of ways.

[0055] According to one embodiment, when a non-traditional abstract link is selected, the resource associated with the link is retrieved, and that resource is displayed by the browser in a manner that replaces the search results page that contained the abstract. In such an embodiment, the selection of the non-traditional abstract link causes a similar operation as selection of most other types of links.

[0056] According to an alternative embodiment, selection of a non-traditional abstract link does not cause the web page containing the abstract to be replaced. Instead, selection of the link causes an HTML overlay to be displayed, where the HTML overlay contains information from or about the resource associated with the link. The overlay itself may contain links which, when selected, cause the corresponding resource to be displayed in place of the web page that contains the abstract.

[0057] The information contained in the overlay may provide the user sufficient information to make it unnecessary to retrieve the entire resource associated with the link. For example, the abstract may be for a hotel's web page, the non-traditional abstract link may be to a page that contains a review of the hotel, and the overlay for the link may contain the first few lines of the review and a rating of the hotel. In many cases, the first few lines of the review and the rating of the hotel will be all that the user desires to see.

[0058] Even if the information contained in the overlay is not sufficient in and of itself, it may provide the user information to help the user decide whether retrieving the resource would be worthwhile. If the user decides, based on the information presented in the overlay, that retrieval of the entire resource associated with the link is desirable, then the user may select the link from within the overlay. As mentioned above, there are many ways in which the link may be selected, and the techniques described herein are not limited to any particular link selection technique.

[0059] An HTML overlay is merely one example of a technique for displaying information associated with non-traditional abstract links without replacing the search results page. Other examples include windows generated by embedded JAVA applets, pop-up windows, etc.

SELECTING WHICH NON-TRADITIONAL ABSTRACT LINKS TO INCLUDE

[0060] In the examples given above, the search engine determines which non-traditional abstract links to include in the abstract of a matching resource. According to one embodiment, the search engine maintains metadata that dictates which non-traditional search links to include in the abstract of a matching resource. This metadata is referred to herein as "link inclusion metadata".

[0061] In one embodiment, the link inclusion metadata maps (1) data that identifies a resource (such as a URL) to (2) data that indicates the non-traditional abstract links to include in the abstract for that resource. In such an embodiment, when a resource matches a search query, the search engine searches the link inclusion metadata based on the URL, and generates an abstract that includes the non-traditional abstract links that are mapped to that URL.

[0062] The link inclusion metadata may be any form of metadata, used by the search engine, to determine which non-traditional abstract links to include in the abstract of a resource. In a simple case, the link inclusion metadata for a resource may expressly identify a link to include in the abstract for the resource. In other cases, the link inclusion metadata may include a fragment of a link, where the search engine generates the entire link based on the fragment of the link and the search query. In yet other cases, the link inclusion metadata may specify rules for generating non-traditional abstract links for a particular resource. The search engine may dynamically generate the non-traditional abstract links, based on the specified rules, at the time the abstract is being generated in response to a specific search query.

CATEGORY-BASED SELECTION OF NON-TRADITIONAL ABSTRACT LINKS

[0063] According to one embodiment, the type of links that a search engine includes in the abstract for a matching resource depends on the nature of the matching resource. For example,

the search engine may embed a link to a “reviews” site within the abstract for the web page of a hotel. However, it does not necessarily make sense to embed a link to a “reviews” site into the abstract for a web page that simply explains how to perform basic arithmetic.

[0064] According to one embodiment, the link inclusion metadata maps resources to categories, and categories to rules about which non-traditional abstract links to include. In such an embodiment, when a resource matches a search query, the search engine determines the source of the resource, and then determines the category to which that source belongs. Based on the category, the search engine determines which rules should dictate what links to embed in the abstract for that matching resource.

[0065] “Auction Sites” is an example of a category. The search engine may store metadata that includes a list of all auction sites. When the page from an auction site matches a search query, the search engine chooses what links to embed in the abstract for the web page based on the fact that the web page is for an auction site. In the auction site example, the matching resource may be a web page for a specific auction of a specific product. Based on the fact that the web page is from an auction site, the search engine may include in the abstract for the web page a link to a web page for all auctions that relate to that type of product. In fact, the search engine may include several links, where each link sorts the auctions for that type of product using a different sort criteria (e.g. ending time, price, etc.)

[0066] The link inclusion metadata may specify a completely different set of rules for different categories of sites. For example, for “Personal Web Page” sites, the search engine may simply embed mid-document links to the terms that matched the search query. For “Company Home Page” sites, the search engine may embed value-added links to map-generating sites. For “Entertainment Source” sites, the search engine may embed value-added links to “review” web sites. These are merely a few examples of how, through metadata that categorizes the sources of searchable resources, the search engine may apply source-specific rules about what links to embed in the abstracts of matching resources.

TWO-PHASE SELECTION OF NON-DEFAULT ITEMS

[0067] According to one embodiment, the selection of which non-default items (including but not limited to non-traditional abstract links) to include in an abstract for a resource is performed in two phases. In the first phase, the search engine identifies a set of “candidate items” based on metadata associated with the resource. In embodiments that make use of non-traditional abstract links, the candidate items may include several non-traditional abstract links.

[0068] In the second phase, a subset of those candidate items is selected based on the search query that matched the resource. Because the subset is selected based on the search query, the abstract for a particular web page may include one set of non-default items when provided as the search results for one query, and may include an entirely different set of non-default items when provided as the search results for a different query.

[0069] In embodiments that make use of non-traditional abstract links, the specific non-traditional abstract links that are included in the abstract for a particular web page will vary based on the search query. In one embodiment, the link inclusion metadata associated with a particular resource may indicate (1) several candidate non-traditional abstract links and (2) selection criteria for each of the candidate non-traditional abstract links. For example, the link inclusion metadata for a web page about electronic devices may include:

- a link to a web page about cell phones, associated with the term “cell phone”;
- a link to another web page about electronic devices; and
- a link to a web page about digital music player, associated with the term “digital music”.

[0070] The search engine selects which of these three candidate non-traditional abstract links to include in the abstract for the matching resource based on the search query. For example, if the search query includes the term “cell phone”, then the search engine may select the link to the web page about cell phones, and the link to the other web page about electronic devices. On the other hand, if the search query includes the term “digital music”, then the search engine may select the link to the web page about digital music players, and the link to the other web page about electronic devices.

[0071] Once the subset of candidate items is selected in this manner, the abstract for the matching item is generated to include only the selected subset of candidate items. Using this technique, the non-traditional abstract links that are included in the abstract for any given matching resource are customized to the likely interests of the search engine user to which they will be presented.

[0072] As another example of the two-phased selection of non-traditional abstract links, assume that the matching resource is a web page that has information about the motorcycle laws for every state in the United States of America. The candidate items may include, for each state, a mid-document link to the information that applies to that specific state. In this example, the selection criteria for each of the mid-document links may be the name of the state. Thus, if the abstract is displayed in response to a search for “California motorcycle laws”, then the abstract will include a mid-document link to the section of the document that

applies to California. On the other hand, if the abstract is displayed in response to a search for "New York motorcycle laws", then the abstract will include a mid-document link to the section of the document that applies to New York.

[0073] FIG. 1 is a flowchart showing the steps involved in the two-phase selection of non-default items. Referring to FIG. 1, at step 100 a query is received by the search engine. At step 102, the search engine identifies matching resources based on the search query. At step 104, the search engine constructs an abstract for each matching resource. The abstract is constructed based both on metadata associated with the resource, and the search query itself.

[0074] Specifically, in one embodiment, step 104 involves locating the metadata that applies to a matching resource (step 106) and determining candidate items based on the metadata (step 108). The candidate items may include many types of non-default items, including one or more non-traditional abstract links. At step 110, a subset of the candidate items are selected based on the search query. At step 112, the search engine generates an abstract that includes the subset of candidate items.

[0075] At step 114, the search engine provides to the user search results that include the abstracts generated in step 104.

FILTERING CRITERIA

[0076] As described above, the second phase of the two-phase selection process selects a subset of the non-default items associated with a matching resource. Thus, the second phase "filters" the candidate items based on some criteria. In the example given above, the criteria relates to the terms used in the search query to which the search engine is responding. However, the search query terms are merely one example of the criteria by which the search engine may filter the candidate items.

[0077] According to one embodiment, the search engine filters the candidate items based, at least in part, on information that the search engine knows about the user that submitted the search query. For example, the search engine may be part of a service for which users register. As part of the registration process, users provide profile information about themselves. When determining which candidate items to filter, the search engine may inspect the profile information for the user that submitted the search query, and filter the candidates based on the information contained therein.

[0078] For example, the profile information for a particular user may indicate that the user has a strong interest in digital photography. A search performed by the user may match a web page that is about several different types of electronic devices. Based on the profile

information for the user, the search engine may limit the non-default items that are included in the abstract for the web page to those that relate to digital photography. The resulting abstract may include, for example, a mid-document link to jump directly to the location within the web page that relates to digital cameras.

[0079] As another example, the profile information for a particular user may indicate that the user resides California. A search performed by the user may match a web page that is about state flags. Based on the profile information for the user, the search engine may limit the non-default items that are included in the abstract for the web page to those that relate to the state flag of California. The resulting abstract may include, for example, a related subject matter link to a different site that focuses more particularly on the state flag of California.

GENERATING LINK INCLUSION METADATA

[0080] As mentioned above, the search engine determines which non-traditional abstract links to include in an abstract based on link inclusion metadata. The link inclusion metadata may come from a variety of sources. For example, in one embodiment, the link inclusion metadata for a resource is generated based on information included in the resource itself. For example, the designer of a web page may include, in the web page itself, information about which non-traditional links to include in the abstract for the hotel's web page. That information, referred to herein as "abstract links hints" may, for example, be demarcated within the web page by specified tags. The search engine may find the abstract links hints based on the tags, and generate the link inclusion metadata for that particular web page based on the abstract links hints. Rather than have a separate abstract links hint section in a resource, the resource may simply use special tags to identify those links, within the resource, that should be promoted to the abstract for the resource.

[0081] In another embodiment, the party that controls a web page may provide link inclusion information directly to the party that controls the search engine. For example, the party that manages a hotel's web page may provide to the party that manages the search engine instructions about which non-traditional abstract links to include in the abstract for the hotel's web page. Base on the information provided by the web page manager, the search engine manager may generate the link inclusion metadata for that particular web page.

[0082] When link inclusion information is provided by the owner of the resource, the link inclusion information may call for more non-traditional abstract links than the search engine provider considers appropriate. Under these conditions, the search engine may be configured

to include no more than a specified maximum number of non-traditional abstract links, even though the link inclusion metadata may call for more.

[0083] In another embodiment, the party that controls the search engine analyzes each resource, and based on the manual analysis, generates the link inclusion metadata for the resource. Because such manual analysis may be more time consuming than other techniques for generating the link inclusion metadata, the use of manual analysis may be limited to only the resources that most frequently appear in search results.

[0084] In yet another embodiment, the search engine is not expressly provided link inclusion information. Instead, the search engine parses a resource, and applies logic to determine the link inclusion metadata for the resource based on the content of the resource. This parsing may be performed, for example, at the time the resource is indexed by the search engine. As mentioned above, the search engine may take into account the category of a resource when generating the link inclusion metadata for the resource. For example, the generation of link inclusion metadata for an auction web site may be based on different logic than the generation of link inclusion metadata for a company web site.

[0085] Information about the category of a resource may be expressly provided to the search engine. The specific link inclusion metadata for the resource is then automatically generated by the search engine based on the logic associated with the specified category, and the contents of the resource itself.

[0086] The category of a resource is merely one example of information that may be used by a search engine in generating the link inclusion metadata for the resource. As another example, the search engine may be provided "link popularity information". Link popularity information is information about how frequently links have been selected by users. Various approaches may be used to obtain link popularity information, and the present techniques are not limited to any particular approach for obtaining link popularity information. Link popularity information may be used by the search engine, for example, to determine which links within a resource are most frequently followed. Based on this determination, the search engine may generate link inclusion information for the resource that identifies the N most popular links within the resource. Based on this link inclusion information, the search engine may generate an abstract for the resource that includes the N most popular links within the resource.

[0087] It should be noted that the various techniques for generating link inclusion metadata that are described herein are not mutually exclusive. Rather, they may be mixed and combined as needed for any given implementation. For example, for any given resource, the link

inclusion metadata may be generated based on (1) information specially demarcated within the resource, (2) link inclusion information provided by the owner of the resource, and (3) logic applied by the search engine to the contents of the resource, based on both link popularity information and resource category information.

ILLUSTRATIVE EXAMPLES

[0088] FIGS. 2-9 illustrate abstracts that may be generated using the techniques described herein. While these examples demonstrate specific embodiments of how the techniques may be implemented, the techniques are not limited to the specific embodiments thus illustrated.

[0089] FIG. 2 illustrates an abstract 200 in which selection of address 208 may activate a value added link that causes a map, associated with the address, to be displayed. Selection of address 212 within the textual body of abstract 202 may cause a similar link to be activated. The telephone number 214 within abstract 202 may be value added information that was retrieved from third-party service, and included in the abstract 202. In the illustrated example, the phone number 214 itself serves as anchor text for a non-traditional abstract link. The non-traditional abstract link associated with phone number 214 may, for example, invoke a third party service that causes the phone number 214 to be sent to the cell phone of the user. Alternatively, if the computer that is displaying the abstract 214 is configured to control a telephone, the non-traditional abstract link associated with phone number 214 may cause the telephone controlled by the computer to call the phone number 214.

[0090] FIG. 3 illustrates abstracts 302 and 304 that may be generated when the matching resources contain structured content. In the example illustrated in FIG. 3, the structured content is from an auction site. Abstract 302 includes information 306, such as the ending date of the particular auction and the immediate purchase price for the particular auction, which has been promoted from the matching document into the abstract. In addition, abstract 302 includes links 308, 310 and 312 which retrieve a related web page (a web page that lists all auctions related to the same item), and causes the related web page to be sorted newly listed first, ending soonest first, and lowest price first, respectively. Abstract 302 also includes a related subject matter link 314 which, when activated, retrieves the web page associated with online stores that sell the same item.

[0091] Abstract 304 may be for the same matching resource as abstract 302, but does not contain the same information. The difference between abstract 304 and abstract 302 may result from the use of query-based selection of candidate items. Thus, because the query that

caused generation of abstract 302 contained different terms than the query that caused generation of abstract 304, the abstracts are not identical.

[0092] Referring to FIG. 4, it illustrates three abstracts 402, 404 and 406 that may be generated using the techniques described herein. Abstract 404 includes a link 408 to a web page associated with the general category 408 to which the item described in the matching resource belongs. Abstract 404 also includes a link 410 to the more specific category to which the item described in the matching resource belongs. Links 404 and 408 may be promoted links from the matching resource itself, or from levels below the matching resource. Links 404 and 408 may have been selected for inclusion in abstract 404 because they had the highest “link popularity” values of all of the links associated with the matching resource.

[0093] Abstract 406 may be for the same matching resource as abstract 404, but does not contain the same information. The difference between abstract 406 and abstract 404 result from the use of query-based selection of candidate items. Thus, because the query that caused generation of abstract 404 contained different terms than the query that caused generation of abstract 406, the abstracts are not identical.

[0094] FIG. 5 illustrates two abstracts 502 and 504 that include various non-traditional abstract links. For example, abstract 502 includes “newly listed”, “ending soonest” and “lowest price first” links which, when activated, retrieve a list of auctions on similar items, and sort the list in the manner indicated by the link.

[0095] FIG. 6 includes six examples of abstracts that may be generated using the techniques described herein. The first example abstract includes direct links to the most popular reviews of the item described in the matching resource. These review links may have been promoted from the matching resource, from levels below the matching resource, or they may simply be related subject matter links that have been associated with the matching resource.

[0096] Example 2 of FIG. 6 includes non-traditional abstract links to various locations within the web site associated with the matching resource. By activating such a link, the user is able to jump directly to a particular location within the matching resource, or to a different web page from the same site as the matching resource.

[0097] Example 3 of FIG. 6 is an example of “deep link extraction”. Specifically, abstract includes a “digital cameras” link 602 that has been promoted from the matching resource. Activation of the digital camera link 602 causes the viewer to jump to a “digital camera web page” one level removed from the matching resource. In addition, the abstract includes links 604 about specific types of digital cameras. Links 604 have been extracted from the digital

camera web page. Thus, activation of one of links 604 causes the viewer to jump two levels removed from the matching resource, to a web page about a specific type of digital camera.

[0098] Example 3 of FIG. 6 is also an example of an abstract generated by the two-phase technique described herein. Specifically, the matching resource is associated with many categories of products (digital cameras, camcorders, televisions, etc.). The link inclusion metadata for the resource may identify links for each of the categories, and links for specific products in each of the categories. The search engine determines which of those links to include in the abstract based on the search query to which the search engine is responding. In the present example, the search query included terms relating specifically to digital cameras. Consequently, the abstract is generated to include links to the web page associated with the digital camera category, and links to the web pages associated with specific types of digital cameras. If the search query had included terms relating to camcorders, then the links relating to digital cameras would not be included in the abstract. Instead, the abstract would include links to a web page associated with camcorders, and to web pages associated with particular types of camcorders.

[0099] Example 4 of FIG. 6 illustrates an abstract that includes a value added service link 606. When link 606 is activated, the link causes the browser to send an address in a request to a third-party mapping service. In response to the request, the third-party service generates a map for the address, and sends the map to the browser.

[0100] Example 5 of FIG. 6 illustrates an abstract that includes links to web sites that are related to the matching document. Specifically, the matching document is a web page about the prices of hotels near Studio Universal. The abstract in Example 5 includes links 608 which, when activated, retrieve web pages associated with specific hotels near Studio Universal. The anchor text for links 608 includes both the name of the specific hotels, and the prices thereof.

[0101] Example 6 of FIG. 6 illustrates an abstract associated with an auction web page. The abstract includes numerous links which, when activated, allow the user to jump immediately to web pages that may be of interest to the user. For example, the abstract shown in Example 6 includes links that allow the user to jump to the web pages associated with specific auctions of the item. The anchor text for those links indicate the time at which each of the corresponding auctions ends.

[0102] FIG. 7 illustrates an abstract associated with the web page of a restaurant. The abstract includes a value added service link which, when selected, submits the address of the restaurant to a mapping service. The mapping service, in turn, generates a map based on the

address and returns the map to the browser. The abstract illustrated in FIG. 7 differs slightly from other mapping service examples, in that the anchor text for the link to the mapping service is “Map”, rather than the address itself.

[0103] FIG. 8 illustrates various ways in which the information associated with a non-traditional abstract link may be presented to a user. Specifically, FIG. 8 illustrates an abstract with a “Weather” link 802. When the Weather link 802 is activated, the link causes the browser to send a request to a weather service site. In response to the request, the weather service site generates weather information, and sends the weather information back to the browser. The browser then displays the weather information in an overlay 804 that does not cause the browser to stop displaying the abstract itself.

[0104] As mentioned above, non-traditional abstract links, such as Weather link 802, may be activated in a variety of ways. For example, Weather link 802 may be activated by hovering the cursor over the “Weather” anchor text, or by clicking on the “Weather” anchor text.

[0105] FIG. 8 also illustrates an abstract that includes a “Map” link 806. When the Map link 806 is activated, a mapping service is called, and the resulting map is presented to the user in an overlay 808.

[0106] Also shown in FIG. 8 is an abstract that includes a value added link 810 which, when activated, causes a URL, address and phone number to be sent to the cell phone of the user. In this example, hovering over the link 810 does not activate the link. Rather, hovering over the link 810 causes the display of an overlay 812 that indicates what will happen if the user activates the link 810.

[0107] FIG. 9 illustrates an example of an abstract that may be generated in response to a search query that specifies a particular hotel, but does not specify a particular location. To assist the user in locating the specified hotel at a specific location, the abstract includes a link 902 which, when activated, causes another search to be performed. However, unlike the original search, the search that will be performed when link 902 is activated searches through resources that are “local” relative to the user. Thus, link 902 is an example of a non-traditional abstract link that is included based, in part, on information that the search engine knows about the user (i.e. the user’s location).

INTERFACES FOR DISPLAYING META INFORMATION

[0108] The interface design illustrated in FIGs. 2-9 should not be considered limiting relative to how metadata about a result may be displayed. For example, any one of a variety of

user interfaces may be used to accommodate display of meta information. For instance, interfaces may be used which allow a user to directly perform the operation on the result site without leaving the search page. For instance, if the matching resource is a search engine page, the search engine might promote the search box into the abstract. As another example, if the result is a music page, then the search engine might promote the play button for a song into the abstract.

[0109] According to one embodiment, the user interfaces change as the user interacts with the page. For instance, if one particular result interests the user, then the search engine might actually change the user interface, and the entire results page, to now "expand" that result to show a rich set of metadata about that result, and the relationship of that result to the user query. In yet another embodiment, the metadata that is displayed is not fixed for a result or a result-query pair. Rather, the metadata may change *for the same result* and *the same query* based on how the user interacts with the page.

ACTIVE ABSTRACTS FOR THINGS OTHER THAN MATCHING RESOURCES

[0110] In examples given above, abstracts are generated by a search engine based on resources that match a search. However, displaying metadata and generating abstracts using the techniques described herein are not limited to the context of search results. For example, the techniques may be applied to define metadata and abstracts for the query itself. A query for instance, can have various senses. The techniques may be employed to provide a *query summarization* that tells the user that jaguar has many senses: cars, cats, operating system. How query-specific metadata is displayed is also an abstracting technique.

[0111] Another use of the active abstract generating techniques described herein involves *result-set* specific summarization. In result-set specific summarization, the abstracts and the metadata change based on the type of results that show up for a query. For instance, if the result set is very diverse, like in the jaguar example, if there are results about the car, cat and operating system all adjacent to each other, then the search engine might choose to show metadata that just highlights the broad category of the result. However, if the result set is uniform – (e.g. many or all of the results are about cats), then the search engine might select metadata that shows detailed in-depth differentiation - one page is a picture of jaguars in Africa (the search engine might choose to the pictures of the cat directly instead of links or words that describe the page), one page talks about biological ancestry of jaguars (the search engine might pull in reference data and biology resources as metadata) and perhaps one page is

about show cats in Las Vegas. Thus, metadata selection criteria can also depend on what the other results look like.

HARDWARE OVERVIEW

[0112] Figure 10 is a block diagram that illustrates a computer system 1000 upon which an embodiment of the invention may be implemented. Computer system 1000 includes a bus 1002 or other communication mechanism for communicating information, and a processor 1004 coupled with bus 1002 for processing information. Computer system 1000 also includes a main memory 1006, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 1002 for storing information and instructions to be executed by processor 1004. Main memory 1006 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 1004. Computer system 1000 further includes a read only memory (ROM) 1008 or other static storage device coupled to bus 1002 for storing static information and instructions for processor 1004. A storage device 1010, such as a magnetic disk or optical disk, is provided and coupled to bus 1002 for storing information and instructions.

[0113] Computer system 1000 may be coupled via bus 1002 to a display 1012, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 1014, including alphanumeric and other keys, is coupled to bus 1002 for communicating information and command selections to processor 1004. Another type of user input device is cursor control 1016, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 1004 and for controlling cursor movement on display 1012. 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.

[0114] The invention is related to the use of computer system 1000 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 1000 in response to processor 1004 executing one or more sequences of one or more instructions contained in main memory 1006. Such instructions may be read into main memory 1006 from another machine-readable medium, such as storage device 1010. Execution of the sequences of instructions contained in main memory 1006 causes processor 1004 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.

[0115] 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 1000, various machine-readable media are involved, for example, in providing instructions to processor 1004 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 1010. Volatile media includes dynamic memory, such as main memory 1006. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 1002. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

[0116] 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.

[0117] Various forms of machine-readable media may be involved in carrying one or more sequences of one or more instructions to processor 1004 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 1000 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 1002. Bus 1002 carries the data to main memory 1006, from which processor 1004 retrieves and executes the instructions. The instructions received by main memory 1006 may optionally be stored on storage device 1010 either before or after execution by processor 1004.

[0118] Computer system 1000 also includes a communication interface 1018 coupled to bus 1002. Communication interface 1018 provides a two-way data communication coupling to a network link 1020 that is connected to a local network 1022. For example, communication interface 1018 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 1018 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 1018 sends and receives

electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0119] Network link 1020 typically provides data communication through one or more networks to other data devices. For example, network link 1020 may provide a connection through local network 1022 to a host computer 1024 or to data equipment operated by an Internet Service Provider (ISP) 1026. ISP 1026 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 1028. Local network 1022 and Internet 1028 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 1020 and through communication interface 1018, which carry the digital data to and from computer system 1000, are exemplary forms of carrier waves transporting the information.

[0120] Computer system 1000 can send messages and receive data, including program code, through the network(s), network link 1020 and communication interface 1018. In the Internet example, a server 1030 might transmit a requested code for an application program through Internet 1028, ISP 1026, local network 1022 and communication interface 1018.

[0121] The received code may be executed by processor 1004 as it is received, and/or stored in storage device 1010, or other non-volatile storage for later execution. In this manner, computer system 1000 may obtain application code in the form of a carrier wave.

[0122] 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.

CLAIMS

What is claimed is:

- 1 1. A method of generating search results in response to a search query, the method
2 comprising:
3 based on the search query, identifying a plurality of matching resources, wherein the
4 plurality of matching resources includes a particular resource;
5 reading metadata associated with said particular resource, wherein said metadata
6 indicates a plurality of candidate items to include in an abstract for said
7 particular resource;
8 based on filtering criteria, selecting a subset of candidate items from said plurality of
9 candidate items; and
10 generating abstracts of each of said plurality of matching resources, wherein the
11 abstract of said particular resource is generated to include the subset of
12 candidate items that were selected based on the search query; and
13 providing said abstracts as search results of said search query.
- 1 2. The method of Claim 1 wherein the step of selecting a subset based on filtering criteria
2 includes selecting a subset based on terms in the search query.
- 1 3. The method of Claim 1 wherein:
2 the method includes reading information about the user that submitted the search
3 query; and
4 the step of selecting a subset based on filtering criteria includes selecting a subset based
5 on the information about the user that submitted the search query.
- 1 4. The method of Claim 1 wherein:
2 said metadata indicates a plurality of candidate non-traditional abstract links to include
3 in an abstract for said particular resource;
4 the step of selecting a subset of candidate items from said plurality of candidate items
5 includes selecting a subset of the candidate non-traditional abstract links from
6 the plurality of candidate non-traditional abstract links; and
7 the step of generating abstracts of each of said plurality of matching resources includes
8 generating an abstract, for said particular resource, that includes the subset of

9 candidate non-traditional abstract links that were selected based on the search
10 query.

1 5. The method of Claim 4 wherein the plurality of candidate non-traditional abstract links
2 includes a mid-document link.

1 6. The method of Claim 4 wherein the plurality of candidate non-traditional abstract links
2 includes a promoted link.

1 7. The method of Claim 6 wherein the promoted link is not contained in, and promoted
2 from at least one level below or above the particular resource.

1 8. The method of Claim 4 wherein the plurality of candidate non-traditional abstract links
2 includes a similar subject matter link that is not contained within the particular resource.

1 9. The method of Claim 4 wherein the plurality of candidate non-traditional abstract links
2 includes a value added service link which, when activated, sends information related to the
3 particular resource to a service that is controlled by a party other than the party that controls
4 the particular resource.

1 10. The method of Claim 9 further comprising the search engine extracting from the
2 particular resource the information related to the particular resource that is sent to said service.

1 11. The method of Claim 9 wherein the service is a service that provides information based
2 on an address, and wherein the information related to the particular resource is an address
3 associated with an entity mentioned in the particular resource.

1 12. The method of Claim 1 wherein the wherein the plurality of candidate items includes a
2 value added information item obtained by the search engine by submitting information
3 associated with the particular resource to a service that is controlled by a party other than the
4 party that controls the particular resource.

1 13. The method of Claim 4 wherein the plurality of candidate non-traditional abstract links
2 includes a related search link which, when activated, initiates a search with terms that are not
3 identical to the terms of said search query.

1 14. The method of Claim 13 further comprising the step of selecting the terms for said
2 search based on both the terms of said search query and terms, not in said search query, that
3 are associated with said particular resource.

1 15. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 1.

1 16. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 2.

1 17. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 3.

1 18. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 4.

1 19. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 5.

1 20. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 6.

1 21. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 7.

1 22. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 8.

1 23. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 9.


1 24. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 10.

1 25. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 11.

1 26. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 12.

1 27. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 13.

1 28. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 14.

**FIG. 1**

2/10




FIG. 2




FIG. 3

4/10




FIG. 4




FIG. 5

6/10




FIG. 6

7/10




FIG. 7

8/10




FIG. 8

9/10





FIG. 9

10/10

FIG. 10



INTERNATIONAL SEARCH REPORT

International application No

PCT/US2005/041074

A. CLASSIFICATION OF SUBJECT MATTER

G06F17/30

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

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

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

EPO-Internal, WPI Data, PAJ, INSPEC, COMPENDEX, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 1 182 581 A (EXALEAD) 27 February 2002 (2002-02-27) abstract paragraph [0001] - paragraph [0027] -----	1-28
X	US 2003/225755 A1 (IWAYAMA MAKOTO ET AL) 4 December 2003 (2003-12-04) abstract -----	1-28
A	YAHOO: "Yahoo Advanced Web Search"[Online] XP002373131 Retrieved from the Internet: URL: http://web.archive.org/web/20040317082739/www.yahoo.com/_y1h=X3oDMTF1M3BpbmVwBF9HA2dsb2JhbF9ncm91cARfUwMyNzE2MTQ5BHR1c3QDMAR0bXBsA25zLWJldGE-/r/so [retrieved on 2004-03-17] the whole document -----	1-28

☐ 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 document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- * & * document member of the same patent family

Date of the actual completion of the international search

23 March 2006

Date of mailing of the international search report

06/04/2006

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

König, W

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

 /US2005/041074

Patent document cited in search report		Publication date		Patent family member(s)	Publication date
EP 1182581	A	27-02-2002	AT	288108 T	15-02-2005
			DE	60017727 D1	03-03-2005
			DE	60017727 T2	29-12-2005
			US	2002052894 A1	02-05-2002
US 2003225755	A1	04-12-2003	JP	2003345810 A	05-12-2003