



(19) **United States**

(12) **Patent Application Publication**
Wen et al.

(10) **Pub. No.: US 2011/0264658 A1**

(43) **Pub. Date: Oct. 27, 2011**

(54) **WEB OBJECT RETRIEVAL BASED ON A LANGUAGE MODEL**

Publication Classification

(75) Inventors: **Ji-Rong Wen**, Beijing (CN);
Shuming Shi, Beijing (CN);
Wei-Ying Ma, Beijing (CN);
Yunxiao Ma, Beijing (CN);
Zaiqing Nie, Beijing (CN)

(51) **Int. Cl.**
G06F 17/30 (2006.01)

(52) **U.S. Cl.** **707/728; 707/E17.014**

(73) Assignee: **Microsoft Corporation**, Redmond, WA (US)

(57) **ABSTRACT**

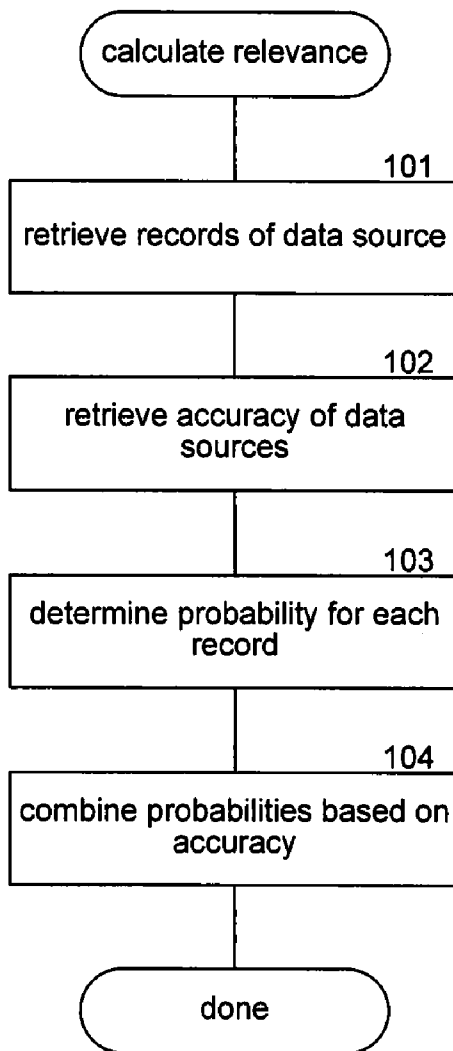
(21) Appl. No.: **13/175,796**

(22) Filed: **Jul. 1, 2011**

A method and system is provided for determining relevance of an object to a term based on a language model. The relevance system provides records extracted from web pages that relate to the object. To determine the relevance of the object to a term, the relevance system first determines, for each record of the object, a probability of generating that term using a language model of the record of that object. The relevance system then calculates the relevance of the object to the term by combining the probabilities. The relevance system may also weight the probabilities based on the accuracy or reliability of the extracted information for each data source.

Related U.S. Application Data

(63) Continuation of application No. 11/459,857, filed on Jul. 25, 2006, now Pat. No. 8,001,130.



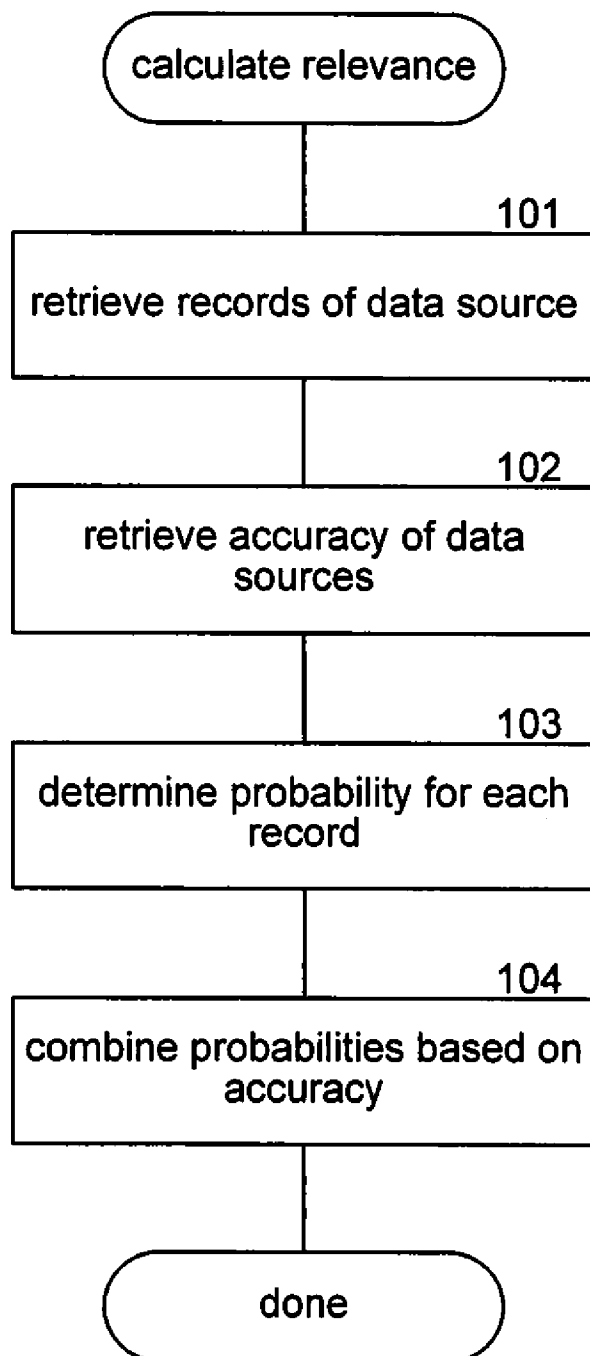


FIG. 1

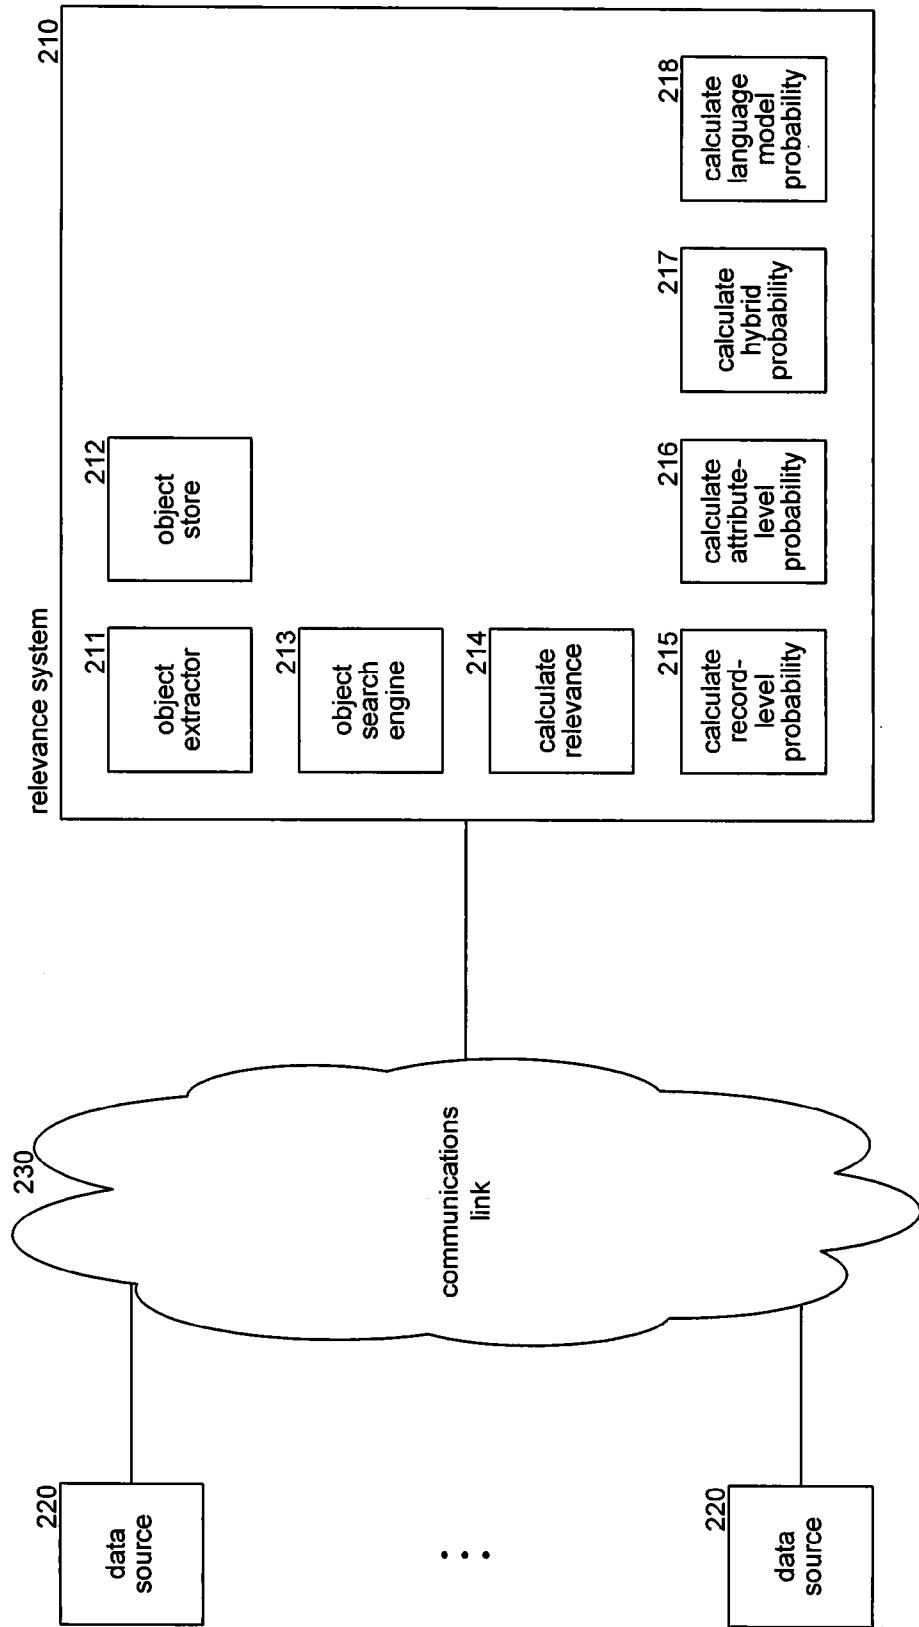


FIG. 2




FIG. 3

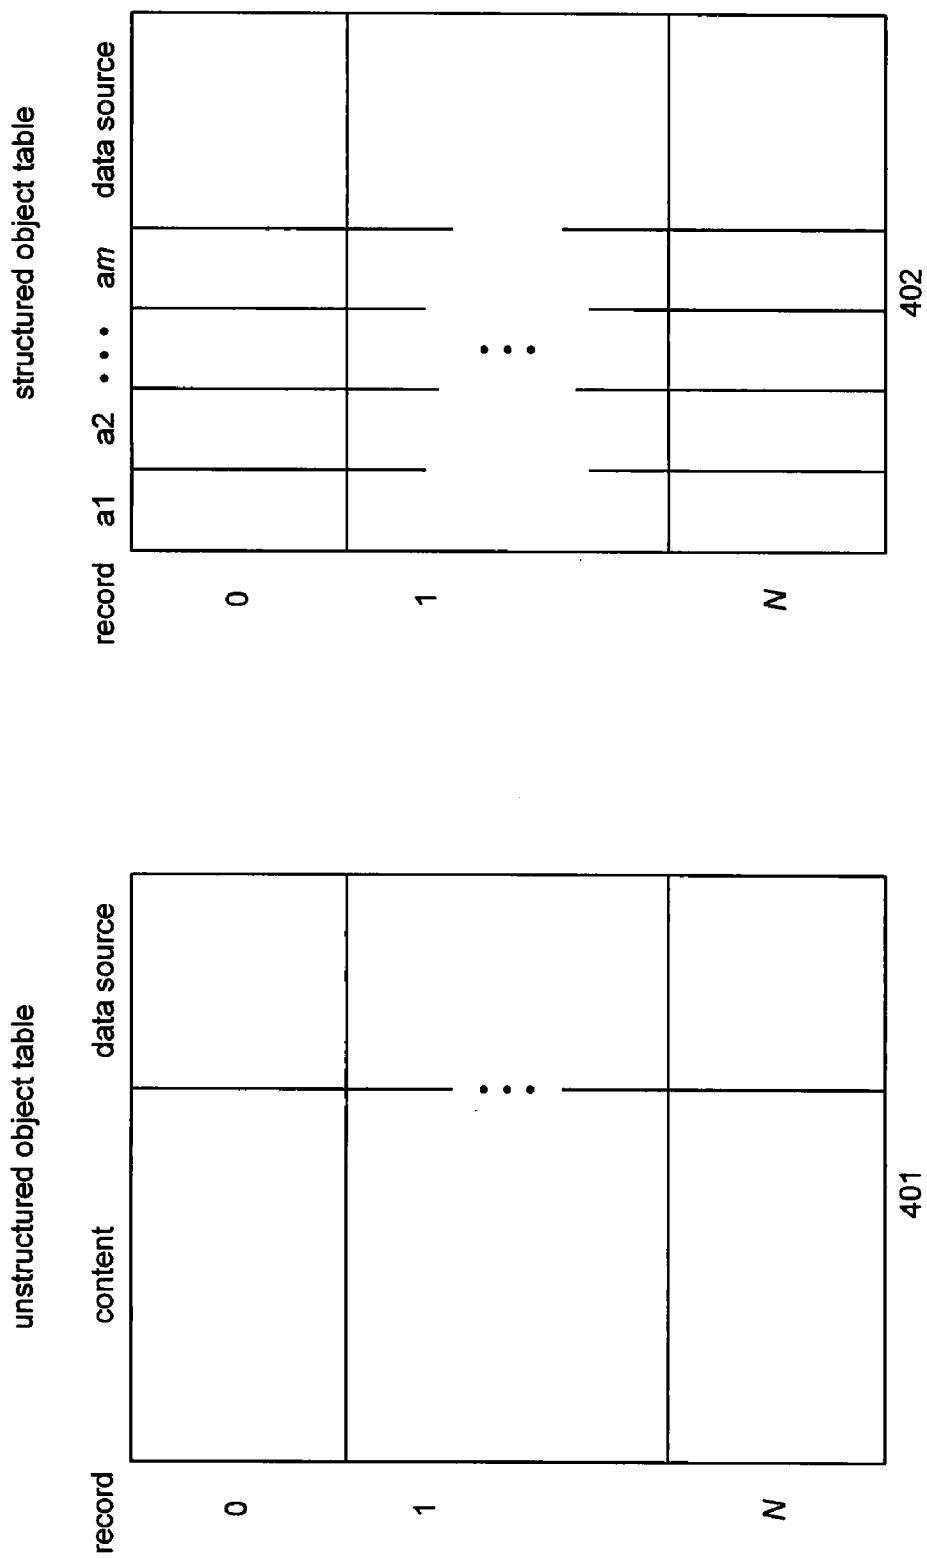


FIG. 4

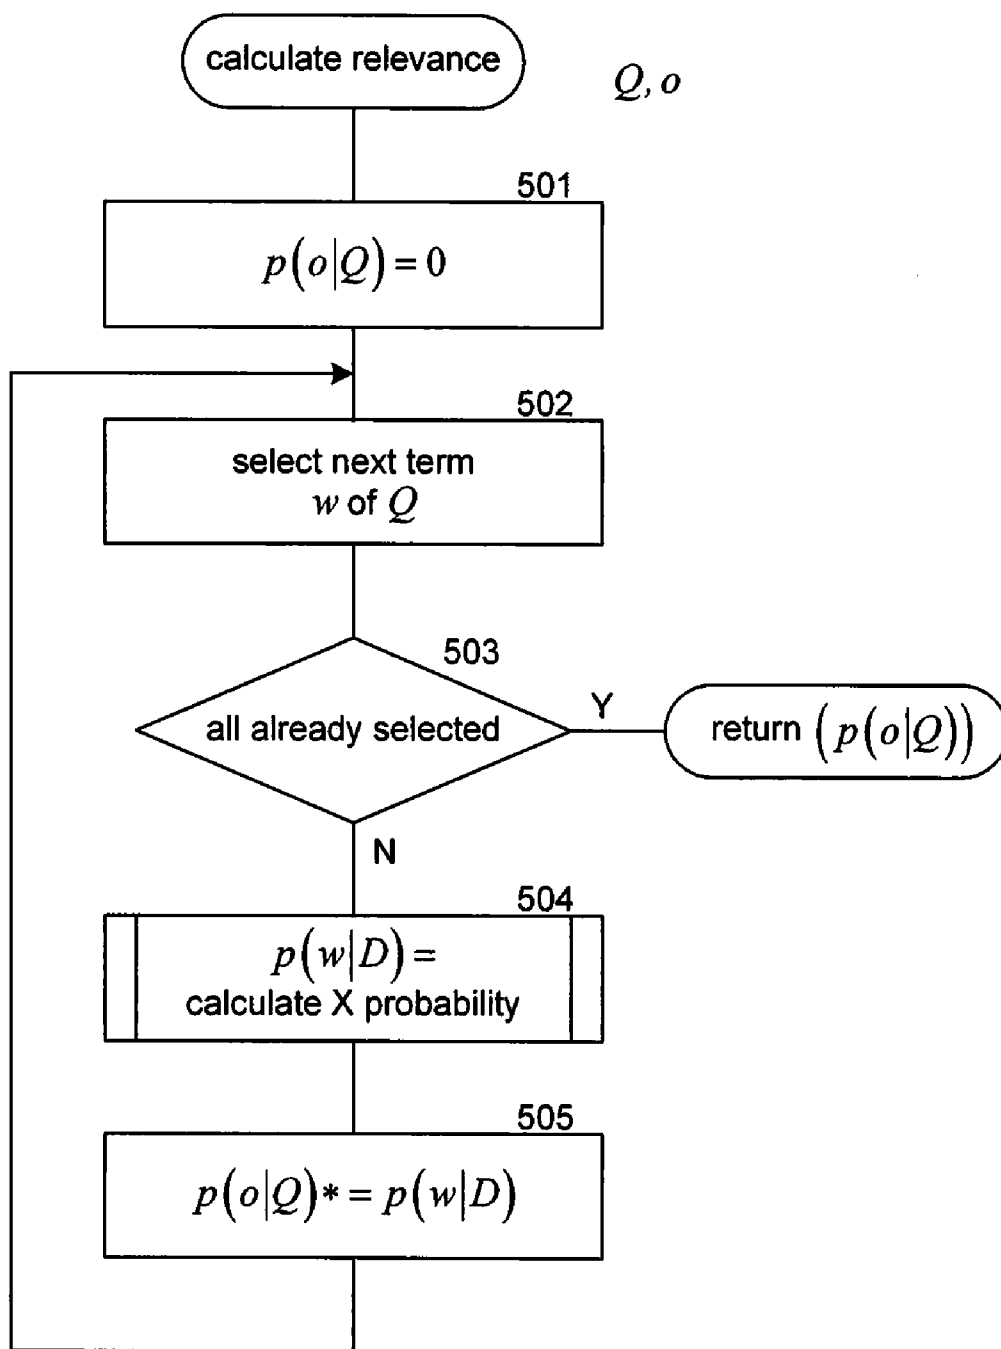


FIG. 5

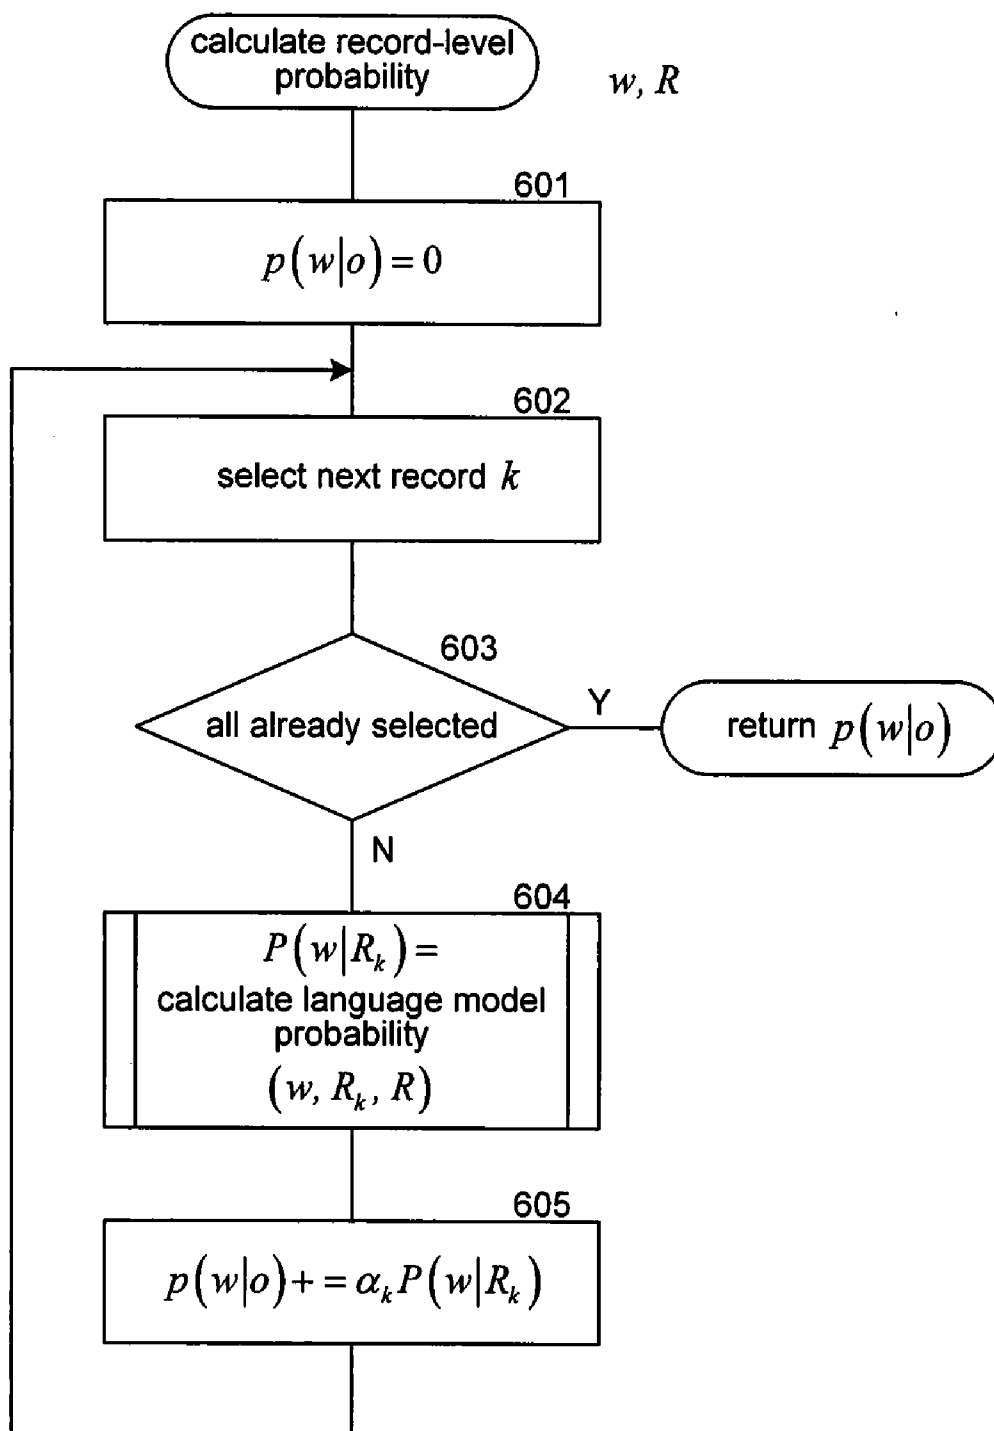


FIG. 6

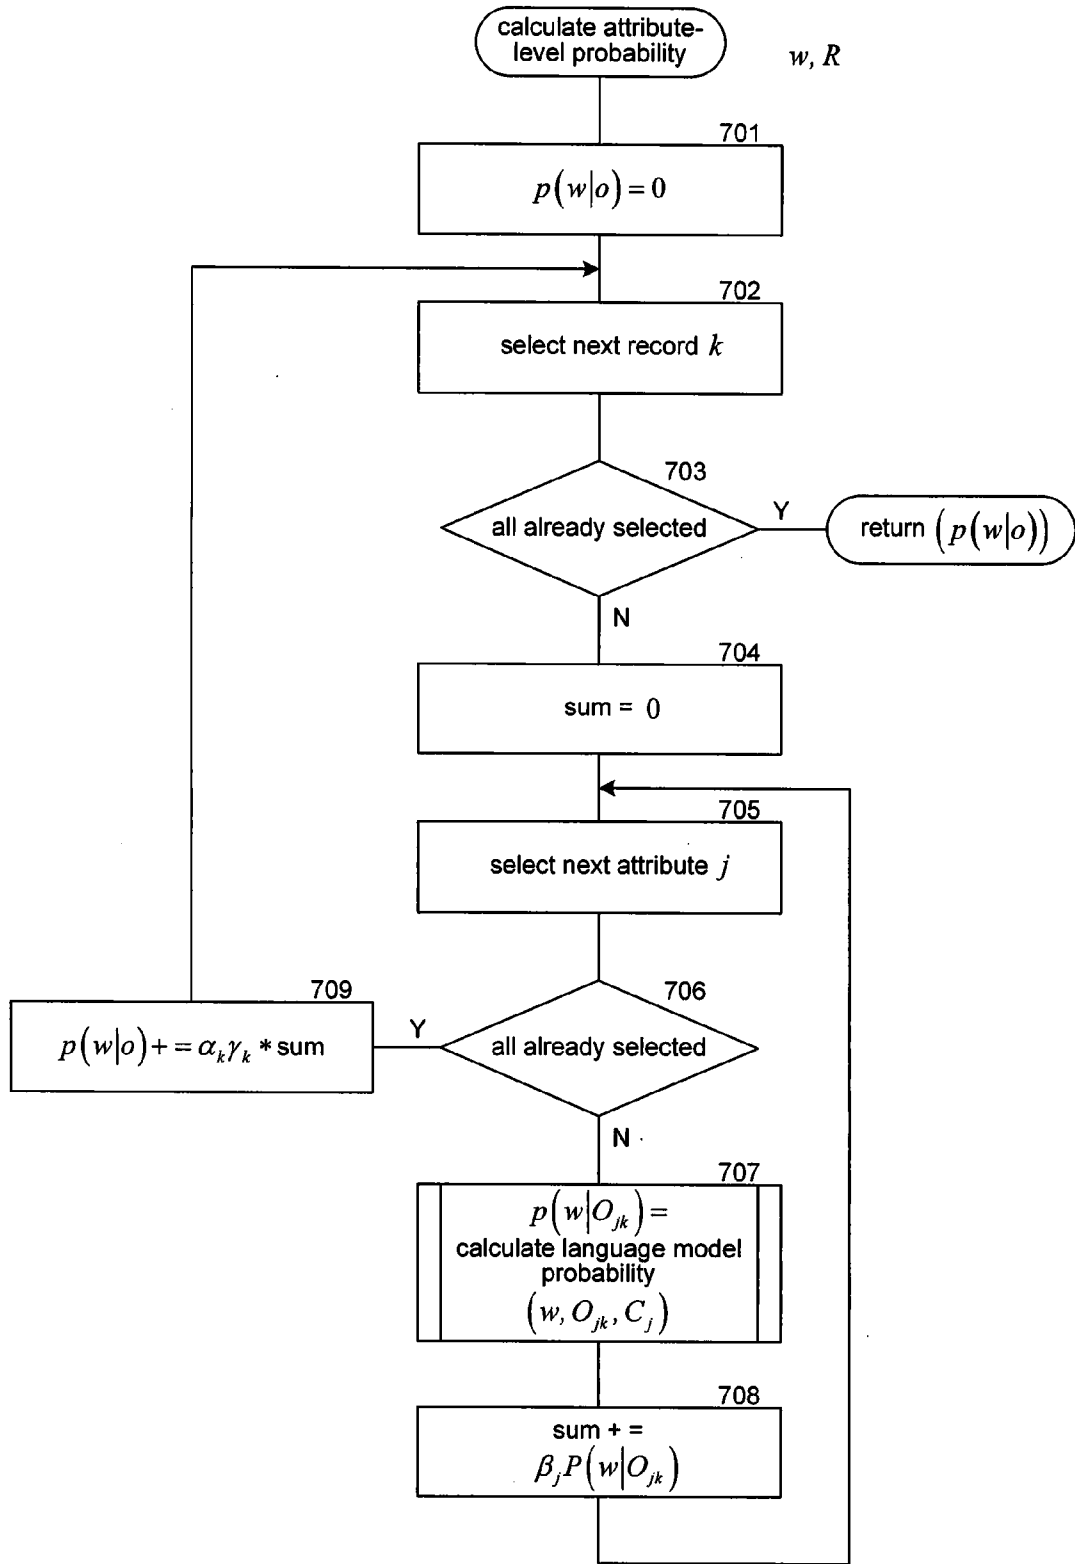


FIG. 7

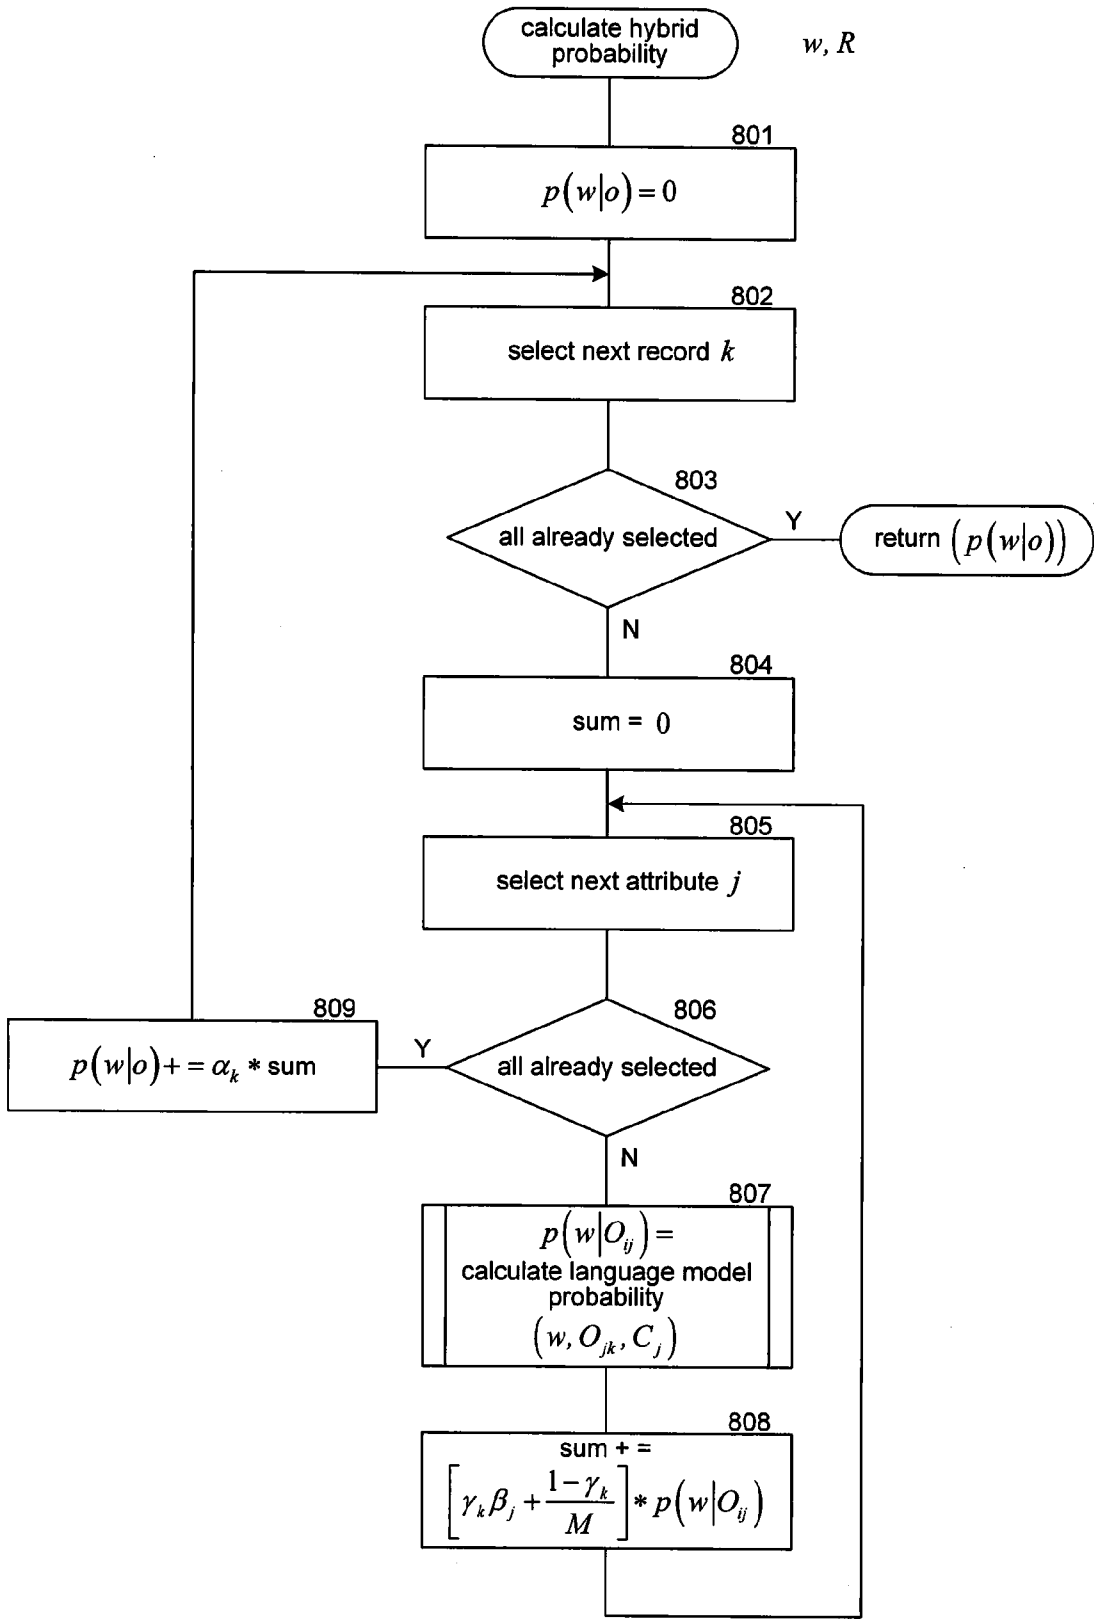


FIG. 8

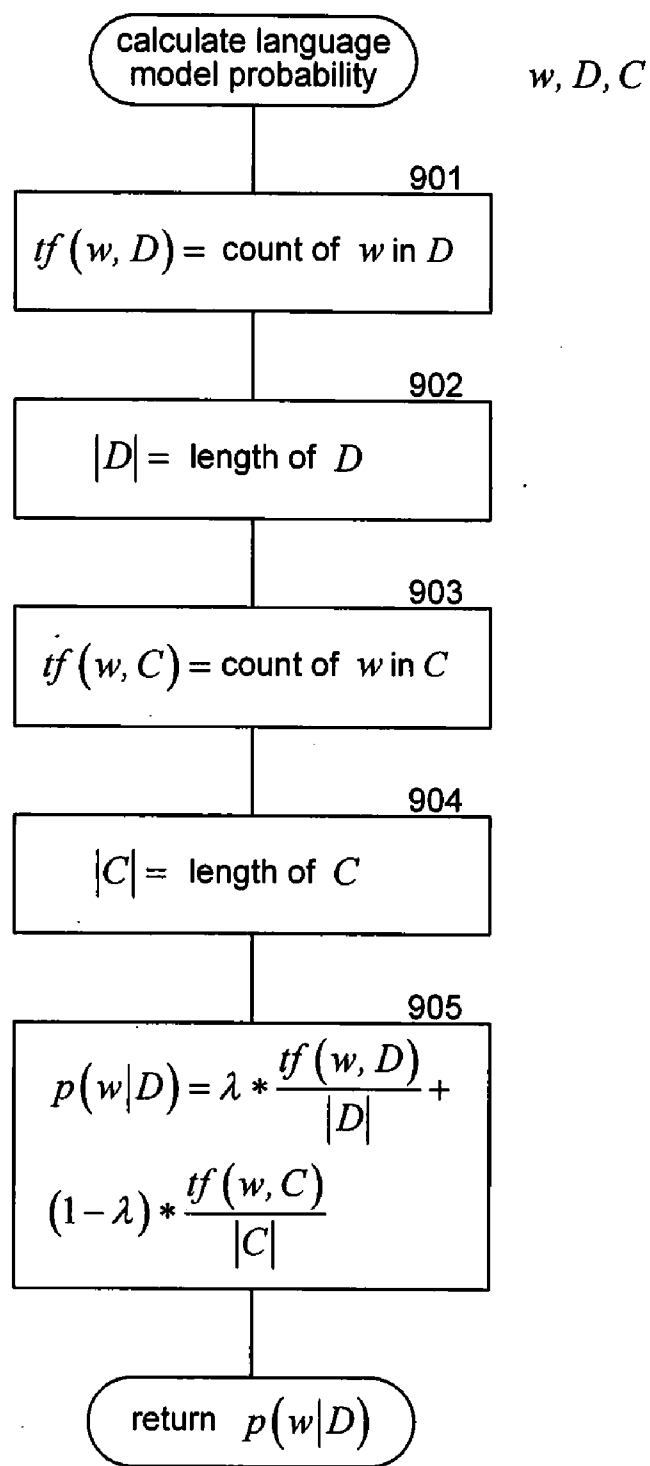


FIG. 9

WEB OBJECT RETRIEVAL BASED ON A LANGUAGE MODEL

CROSS-REFERENCE TO RELATED APPLICATION

[0001] The present application is a Continuation of U.S. patent application Ser. No. 11/459,857, filed Jul. 25, 2006, and entitled “WEB OBJECT RETRIEVAL BASED ON A LANGUAGE MODEL,” which is incorporated herein in its entirety by reference.

BACKGROUND

[0002] Many search engine services, such as Google and Overture, provide for searching for information that is accessible via the Internet. These search engine services allow users to search for display pages, such as web pages, that may be of interest to users. After a user submits a search request (also referred to as a “query”) that includes search terms, the search engine service identifies web pages that may be related to those search terms. To quickly identify related web pages, a search engine service may maintain a mapping of keywords to web pages. The search engine service may generate this mapping by “crawling” the web (i.e., the World Wide Web) to extract the keywords of each web page. To crawl the web, a search engine service may use a list of root web pages and identify all web pages that are accessible through those root web pages. The keywords of any particular web page can be extracted using various well-known information retrieval techniques, such as identifying the words of a headline, the words supplied in the metadata of the web page, the words that are highlighted, and so on. The search engine service may calculate a score that indicates how to rank the web pages based on the relevance of each web page to the search request, web page popularity (e.g., Google’s PageRank), and so on. The search engine service then displays to the user the links to those web pages in the order indicated by the scores.

[0003] A web page may contain information about various types of objects such as products, people, papers, organizations, and so on, which are referred to as “web objects.” For example, one web page may contain a product review of a certain model of camera, and another web page may contain an advertisement offering to sell that model of camera at a certain price. As another example, one web page may contain a journal article, and another web page may be the homepage of an author of the journal article. A person who is searching for information about an object may need information that is contained in different web pages. For example, a person who is interested in purchasing a certain camera may want to read reviews of the camera and to determine who is offering the camera at the lowest price.

[0004] To obtain such information, a person would typically use a search engine to find web pages that contain information about the camera. The person would enter a search query that may include the manufacturer and model number of the camera. The search engine then identifies web pages that match the search query and presents those web pages to the user in an order that is based on how relevant the content of each web page is to the search query. The person would then need to view the various web pages to find the desired information. For example, the person may first try to find web pages that contain reviews of the camera. After

reading the reviews, the person may then try to locate a web page that contains an advertisement for the camera at the lowest price.

[0005] To make it easier to access information about web objects, many systems have been developed to extract information about web objects from web pages. Web pages often allocate a record for each object that is to be displayed. For example, a web page that lists several cameras for sale may include a record for each camera. Each record contains attributes of the object such as an image of the camera, its make and model, and its price. The extraction of such information can be difficult because web pages contain a wide variety of layouts of records and layouts of attributes within records.

[0006] Users can submit queries to a search system to locate information about web objects of interest in a manner similar to how users submit queries to locate web pages of interest. When a user submits a query to locate web object information, traditional database-type retrieval techniques can be used to search for a web object with attributes that match the query. These traditional techniques when applied to web objects are not particularly effective because they assume that the underlying data is reliable. The extraction of web object information can, however, be unreliable for several reasons. First, it can be difficult to precisely identify the record or the portion of a web page that corresponds to a web object. For example, it can be difficult to determine whether adjacent text represents data for the same object or two different objects. If a record is not identified correctly, then the identification of the attributes will likely not be correct. Second, even if the identification of a record is correct, the attributes of the record may still be incorrectly identified. For example, it can be difficult to determine whether a certain number in a record corresponds to the weight of the product, a dimension of the product, and so on. Third, the data source that provides the web page may itself provide unreliable data. For example, a web page advertising a product may simply have wrong information such as the wrong manufacturer for a certain model number of television. Because of this unreliability of extracted web object information, systems that perform searches based on extracted information on web objects often do not provide satisfactory results.

SUMMARY

[0007] A method and system is provided for determining relevance of an object to a term based on a language model. The relevance system provides records extracted from web pages that relate to the object. To determine the relevance of the object to a term, the relevance system first determines, for each record of the object, a probability of generating that term using a language model of the record of that object. The relevance system then calculates the relevance of the object to the term by combining the probabilities. The relevance system may also weight the probabilities based on the accuracy or reliability of the extracted information for each data source. A search system can use the relevance system to calculate relevance of an object to a query by calculating the relevance of the object to each term of the query and then combining the relevances for each term into a relevance for the query.

[0008] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the

claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] FIG. 1 is a flow diagram that illustrates high-level processing of a calculate relevance component of the relevance system in one embodiment.

[0010] FIG. 2 is a block diagram that illustrates components of the relevance system in one embodiment.

[0011] FIG. 3 is a block diagram that illustrates a data structure of the object store in one embodiment.

[0012] FIG. 4 is a block diagram that illustrates a structured and an unstructured object table in one embodiment.

[0013] FIG. 5 is a flow diagram that illustrates low-level processing of the calculate relevance component of the relevance system in one embodiment.

[0014] FIG. 6 is a flow diagram that illustrates the processing of the calculate record-level probability component in one embodiment.

[0015] FIG. 7 is a flow diagram that illustrates the processing of the calculate attribute-level probability component of the relevance system in one embodiment.

[0016] FIG. 8 is a flow diagram that illustrates the processing of the calculate hybrid probability component of the relevance system in one embodiment.

[0017] FIG. 9 is a flow diagram that illustrates the processing of the calculate language model probability component of the relevance system in one embodiment.

DETAILED DESCRIPTION

[0018] A method and system for determining relevance of an object to a term based on a language model factoring in accuracy of extracted information relating to the object is provided. In one embodiment, the relevance system provides records extracted from web pages that relate to the object. For example, the relevance system may extract a record from each of 10 web pages for a particular make and model of a television. The relevance system may also identify the attributes of the records. The collection of all the terms of a record may be considered to be an unstructured representation of the object, and the attributes of the record may be considered to be a structured representation of the object. To determine the relevance of the object to a term, the relevance system first determines, for each record of the object, a probability of generating that term using a language model of the record of that object. A language model provides a probability distribution of terms or sequences of terms occurring in documents that conform to the language model. The relevance system considers each record of an object to be a document and the collection of records for multiple (e.g., all) objects to be a corpus of documents relating to that object. The relevance system also provides metrics indicating the accuracy or reliability of the extracted information for each data source. The relevance system may provide a record accuracy metric that rates how accurately the records of a certain data source can be identified. The relevance system may also provide an attribute accuracy metric that rates how accurately the attributes of a certain data source can be identified. The record accuracy metric and the attribute accuracy metric may both factor in a data accuracy metric, which indicates reliability of the data provided by the data source. The relevance system then calculates the relevance of the object to the term by combining the probabilities of the records, factoring in the

accuracy of the data sources of the records. For example, an accuracy metric may range from 0 to 1 with 1 indicating a high accuracy. The relevance system may then add the probabilities of the records weighted by the accuracy metric of the data source of the record. Thus, a record from an inaccurate data source will have very little effect on the relevance of the object, whereas a record from an accurate data source will have a significant effect on the relevance of the object. A search system can use the relevance system to calculate relevance of an object to a query by calculating the relevance of the object to each term of the query and then combining the relevances for each term into a relevance for the query. In this way, the relevance system calculates the relevance of an object to a term by factoring in the accuracy of the extracted information relating to the object. The relevance system may also determine relevance of an object to a query based on a language model without factoring in the accuracy of the extracted information. Conversely, the relevance system may determine relevance of an object to a query factoring in the accuracy of the extracted information without using a language model.

[0019] In one embodiment, the relevance system calculates the probability of an object being relevant to a term based on an analysis of each record as a collection of terms (record-level probability), of each attribute as a collection of terms (attribute-level probability), or of a combination of records as a collection of terms and attributes as a collection of terms (hybrid probability). The relevance system calculates the record-level probability for each record of an object using a language model that treats each record as a document that contains the terms of the record and the corpus of documents including the records of multiple objects. The relevance system then weights the probability of each record by the record accuracy of the data source of the record and adds the weighted probabilities together.

[0020] The relevance system calculates the attribute-level probability for each record using a different language model for each attribute of the object. For example, if an object representing a product has the attributes of product description and product manufacturer, then the relevance system will have one language model for the product description attribute and another for the product manufacturer attribute. The relevance system treats the attribute of each record as a document that contains the terms of that attribute and the corpus of documents including the terms of the attributes of records of multiple objects. For example, if an object has a product manufacturer attribute and 10 records have been extracted, the relevance system considers the terms of the manufacturer attribute of each record to be a separate document, resulting in 10 separate documents. The relevance system then uses the language model of each attribute to calculate a probability of generating that term using a language model for the attribute of each record of the object. Thus, the relevance system generates a probability for each attribute of each record. The relevance system then combines the probabilities of the attributes for each record to give a probability for that record. For example, if an object has the attributes of product description, product manufacturer, and product model number, the relevance system calculates an attribute probability for each attribute for a record and adds them together to give an overall record probability for the record being relevant to the term. The relevance system may weight the attribute probabilities based on how important the attribute is to the information content of the record. For example, a model number attribute

may be very important because an exact match between a term and the model number may indicate a high degree of relevance, whereas a color attribute may not be very important because an exact match between a term and the color may not be particularly relevant. The relevance system then combines the record probabilities to give a probability of the object being relevant to the term. The relevance system may weight the record probabilities based on the accuracy of the data source of the record before combining the probabilities.

[0021] The relevance system calculates the hybrid probability for each record using a different language model for each attribute of the object and factoring in the importance of each attribute based on the accuracy of the data source. The relevance system calculates attribute probabilities as described above for attribute-level probabilities. The relevance system, however, weights the attribute probabilities based on importance to the extent that the data source accurately identifies attributes. If a data source has a high attribute accuracy, then the relevance system weights the attribute probabilities by the importance. If a data source, however, has a low attribute accuracy, the relevance system weights the attributes equally under the assumption that one inaccurate attribute should not be weighted more heavily than another inaccurate attribute. If a data source has a medium attribute accuracy, the relevance system weights the attributes somewhere between the importance of the attribute and an equal importance for all the attributes of the record.

[0022] The relevance system may use accuracy information related to a data source that is based on a person's objective assessment of the accuracy. To determine the record accuracy, people may review web pages of a data source and records extracted from web pages and enter what they think is an appropriate value for the record accuracy metric for that data source. The metrics may then be averaged or in some other way combined to give a final metric. The attribute accuracy and the data accuracy of the data source may be determined in a similar manner. Alternatively, the relevance system may use an automated accuracy assessment. For example, a classifier may be trained to automatically assess the accuracy of the data source.

[0023] Language models typically determine the relevance of a document to a query based on the probability of generating the query from the language model of the document, as represented by the following equation:

$$P(D|Q) \propto P(Q|D) \cdot P(D) \tag{1}$$

where D represents the document, Q represents the query, and P(A|B) represents the probability of A given B. If the terms of the query are assumed to be independent, then the relevance system represents the probability of generating a query from the language model of the document by the following equation:

$$P(Q|D) = \prod_{i=1}^{|Q|} P(w_i|D) \tag{2}$$

where w_i represents the i th term of Q, $|Q|$ represents the length of Q, and $P(w_i|D)$ represents the probability of generating term w_i from the language model of D.

[0024] The relevance system represents the probability of generating a term from a language model of a document by the following equation:

$$P(w|D) = \lambda \cdot \frac{tf(w, D)}{|D|} + (1 - \lambda) \cdot \frac{tf(w, C)}{|C|} \tag{3}$$

where $|D|$ represents the length of document D, $tf(w,D)$ represents the term frequency (i.e., number of occurrences) of term w in D, $|C|$ represents the number of occurrences of terms in the whole collection, and $tf(w,C)$ represents the term frequency of term w in the whole collection C. The parameter λ represents a Dirichlet smoothing and can have a value between 0 and 1. The relevance system may let the parameter λ rely on document length $|D|$ as represented by the following equation:

$$\lambda = \frac{|D|}{|D| + \mu} \tag{4}$$

where μ represents a parameter that may be set based on the average document length in the collection.

[0025] The relevance system calculates the relevance of an object to a term based on record-level probabilities according to the following equation:

$$p(w|o) = \sum_{k=1}^K (\alpha_k P(w|R_k)) \tag{5}$$

where o represents the object, $P(w|R_k)$ represents the probability of generating w from the record R_k , K is the number of records for the object o and α_k is the record accuracy such that

$$\sum_k \alpha_k = 1.$$

The relevance system calculates $P(w|R_k)$ by treating each record R_k as a document. Therefore, by using Equation 3, the relevance system represents $P(w|R_k)$ by the following equation:

$$P(w|R_k) = \lambda \frac{tf(w, R_k)}{|R_k|} + (1 - \lambda) \cdot \frac{tf(w, C)}{|C|} \tag{6}$$

where C is the collection of all the records for all the objects.

[0026] The relevance system calculates the relevance of an object to a term based on attribute-level probabilities according to the following equation:

$$P(w|O) = \sum_{k=1}^K \left(\alpha_k \gamma_k \sum_{j=1}^M \beta_j P(w|O_{jk}) \right) \tag{7}$$

where γ_k represents the attribute accuracy, $\alpha_k \gamma_k$ represents a normalized accuracy of both the record accuracy and attribute accuracy of record k such that

$$\sum_k \alpha_k \gamma_k = 1,$$

represents the number of attributes, represents the importance of the attribute j such that

$$\sum_j \beta_j = 1,$$

O_{jk} represents the value of the attribute j of the record k, and $P(w|O_{jk})$ represents the probability of generating w from the attribute j of record k. The relevance system calculates $P(w|O_{jk})$ treating each O_{jk} as a document. Therefore, by using Equation 3, the relevance system represents $P(w|O_{jk})$ by the following equation:

$$P(w|O_{jk}) = \lambda \cdot \frac{f(w, O_{jk})}{|O_{jk}|} + (1 - \lambda) \cdot \frac{f(w, C_j)}{|C_j|} \tag{8}$$

where C_j represents the collection of all the attributes j of all records of all objects in a collection of objects.

[0027] The relevance system calculates the relevance of an object to a term based on hybrid probabilities according to the following equation:

$$P(w|O) = \sum_{k=1}^K \left(\alpha_k \sum_{j=1}^M \left(\gamma_k \beta_j + (1 - \gamma_k) \frac{1}{M} \right) P(w|O_{jk}) \right) \tag{9}$$

The relevance system thus weights $P(w|O_{jk})$ between the importance of attribute j and an equal weighting based on the attribute accuracy.

[0028] FIG. 1 is a flow diagram that illustrates high-level processing of a calculate relevance component of the relevance system in one embodiment. The component calculates the relevance of an object to a term. In block 101, the component retrieves the records of various data sources for various objects. In block 102, the component retrieves the accuracy of the data sources. The accuracies may include a record accuracy and an attribute accuracy that each factor in the data accuracy of the data source. In block 103, the component determines for each record of the object the probability that the record is relevant to the term. In block 104, the component combines the probabilities to give an overall probability that the object is relevant to the term. The component may weight the probabilities when combining them based on the accuracies of their data sources.

[0029] FIG. 2 is a block diagram that illustrates components of the relevance system in one embodiment. The relevance system 210 may be connected to various data sources 220 via communications link 230. The data sources may be web sites that serve web pages that contain web objects, file servers that contain documents having objects, and so on. The relevance system may include an object extractor component 211 and an object store 212. The object extractor component may crawl the data sources and extract object information as

records and attributes within records. The object extractor component may identify objects, group record information for the same object, and store the record information in the object store. The relevance system may also include an object search engine 213, a calculate relevance component 214, a calculate record-level probability component 215, a calculate attribute-level probability component 216, a calculate hybrid probability component 217, and a calculate language model probability component 218. The object extractor component and object search engine may be implemented on separate computing system and not be part of the relevance system. The object search engine may receive a query from a user and search the object store for objects that are related to the query. To determine the relevance of an object to the query, the object search engine invokes the calculate relevance component. The object search engine then combines the relevances to give an overall probability that the object is related to the query. The relevance system may use either the calculate record-level probability component, the calculate attribute-level probability component, or the calculate hybrid probability component for calculating the probability that the object is relevant to a term. The calculate language model probability component implements the language model for a document (e.g., record or attribute). The relevance system may also include a store that contains the record accuracy, attribute accuracy, and data accuracy of the various data sources.

[0030] FIG. 3 is a block diagram that illustrates a data structure of the object store in one embodiment. The data structure is an object table 301 that contains an entry for each object. Each entry includes an identifier of the object plus a reference to a structured object table 303 or unstructured object table 302 depending on whether the information extracted from data sources for the object is structured or unstructured. The extractor component populates the object table as it crawls the data sources. FIG. 4 is a block diagram that illustrates a structured and an unstructured object table in one embodiment. An unstructured object table 401 contains an entry for each record identified for the object. Each entry contains the content of the record and an indication of the data source. The content of the record includes all the terms (e.g., words and numbers) from the portion of the document (e.g., web page) that is identified as the record. A structured object table 402 contains an entry for each record identified for the object. Each entry contains a field for each attribute and an indication of the data source. Each field contains the value of the corresponding attribute that is identified from the portion of the document that is identified as the record. These data structures represent a logical organization of the data. The data structures may be implemented using various well-known techniques such as link lists, hash tables, and so on.

[0031] The computing devices on which the relevance system may be implemented may include a central processing unit, memory, input devices (e.g., keyboard and pointing devices), output devices (e.g., display devices), and storage devices (e.g., disk drives). The memory and storage devices are computer-readable media that may contain instructions that implement the relevance system. In addition, the data structures and message structures may be stored or transmitted via a data transmission medium, such as a signal on a communications link. Various communications links may be used, such as the Internet, a local area network, a wide area network, or a point-to-point dial-up connection.

[0032] The relevance system may be used to monitor configuration in various operating environments that include per-

sonal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

[0033] The relevance system may be described in the general context of computer-executable instructions, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, objects, components, data structures, and so on that perform particular tasks or implement particular abstract data types. Typically, the functionality of the program modules may be combined or distributed as desired in various embodiments.

[0034] FIG. 5 is a flow diagram that illustrates low-level processing of the calculate relevance component of the relevance system in one embodiment. The component is passed a query and an object and calculates the relevance of the object to the query. In block 501, the component initializes to zero a probability that a language model for the object will generate the query. In blocks 502-505, the component loops calculating a probability for each term within the query and accumulating the probability for the query. In block 502, the component selects the next term of the query. In decision block 503, if all the terms have already been selected, then the component returns the probability of the query, else the component continues at block 504. In block 504, the component invokes a calculate probability component to calculate the probability that the language model for the document will generate the term. The component may invoke the record-level component, the attribute-level component, or the hybrid component. In block 505, the component accumulates the probability for the term into a probability for the query and then loops to block 502 to select the next term.

[0035] FIG. 6 is a flow diagram that illustrates the processing of the calculate record-level probability component in one embodiment. The component is passed a term and an object and calculates the probability that the object is relevant to the term. In block 601, the component initializes a probability of relevance to zero. In blocks 602-605, the component loops accumulating the probability for each record of the object. In block 602, the component selects the next record of the object. In decision block 603, if all the records have already been selected, then the component returns the probability, else the component continues at block 604. In block 604, the component invokes a calculate language model probability component to calculate a record-level probability that the language model for the record will generate the term. In block 605, the component accumulates the record level probabilities weighted by the record accuracy of the data source into a probability for the object and then loops to block 602 to select the next record.

[0036] FIG. 7 is a flow diagram that illustrates the processing of the calculate attribute-level probability component of the relevance system in one embodiment. The component is passed a term and an object and calculates the probability that the object is relevant to the term. In block 701, the component initializes the probability to zero. In blocks 702-709, the component loops selecting each record for the object and each attribute within each record to accumulate the probability for the object. In block 702, the component selects the next record of the object. In decision block 703, if all the records have already been selected, then the component returns the

probability of the object, else the component continues at block 704. In block 704, the component initializes a sum of the contribution of the probabilities of the attributes of the selected record to zero. In block 705, the component selects the next attribute of the selected record. In decision block 706, if all the attributes of the selected record have already been selected, then the component continues at block 709, else the component continues at block 707. In block 707, the component invokes the calculate language model probability component to calculate the probability that a language model of the attribute will generate the term. In block 708, the component aggregates the probabilities of the attributes weighted by the importance of the attributes and then loops to block 705 to select the next attribute. In block 709, the component aggregates the sum of the attributes weighted by the record accuracy and attribute accuracy into the probability for the object and then loops to block 702 to select the next record.

[0037] FIG. 8 is a flow diagram that illustrates the processing of the calculate hybrid probability component of the relevance system in one embodiment. The component is passed a term and an object and calculates the probability that the object is relevant to the term. In block 801, the component initializes the probability to zero. In blocks 802-809, the component loops selecting each record for the object and each attribute within each record to accumulate the probability for the object. In block 802, the component selects the next record of the object. In decision block 803, if all the records have already been selected, then the component returns the probability of the object, else the component continues at block 804. In block 804, the component initializes a sum of the contribution of the probabilities of the attributes of the selected record to zero. In block 805, the component selects the next attribute of the selected record. In decision block 806, if all the attributes of the selected record have already been selected, then the component continues at block 809, else the component continues at block 807. In block 807, the component invokes the calculate language model probability component to calculate the probability that a language model of the attribute will generate the term. In block 808, the component aggregates the probabilities of the attributes weighted by the importance of the attributes adjusted by the attribute accuracy and then loops to block 805 to select the next attribute. In block 809, the component aggregates the sum of the attributes weighted by the record accuracy into the probability for the object and then loops to block 802 to select the next record.

[0038] FIG. 9 is a flow diagram that illustrates the processing of the calculate language model probability component of the relevance system in one embodiment. The component is passed a term, a document, and a collection of documents and determines the probability that a language model of the document will generate the term. In block 901, the component counts the number of occurrences of the term within the document. In block 902, the component counts the number of occurrences of all terms within the document. In block 903, the component counts the number of occurrences of the term within the collection of documents. In block 904, the component counts the number of occurrences of all terms within the collection of documents. In block 905, the component calculates the probability using Equation 3 and then completes.

[0039] Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific

features and acts described above are disclosed as example forms of implementing the claims. For example, a system may extract object information from web pages using conventional techniques combined with factoring in the accuracy of the data source. The conventional techniques may include data record extraction as described in Liu, B., Grossman, R., and Zhai, Y., "Mining Data Records in Web Pages," ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2003; Wang, J. and Lochovsky, F. H., "Data Extraction and Label Assignment for Web Databases," World Wide Web Conference, 2003; Lerman, K., Getoor, L., Minton, S., and Knoblock, C. A., "Using the Structure of Web Sites for Automatic Segmentation of Tables," ACM SIGMOD Conference, 2004; Tejada, S., Minton, C. A., and Knoblock, A., "Learning Domain-Independent String Transformation Weights for High Accuracy Object Identification," ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2002; and Zhu, J. Nie, Z., Wen, J., Zhang, B., and Ma, W., "2D Conditional Random Fields for Web Information Extraction," Proceedings of the 22nd International Conference on Machine Learning, 2005. These conventional techniques can weight the contribution of each data source according to its accuracy. Also, when determining whether an object is relevant to a query, the relevance system may use a conventional technique for determining a likelihood that the object is relevant to the query augmented by factoring in the accuracy of data sources of records of the object. Accordingly, the invention is not limited except as by the appended claims.

1-20. (canceled)

21. A computer-readable storage medium containing instructions for controlling a computer system with a processor and a memory to determine relevance of an object to a target term, by a method comprising:

generating a collection of a plurality of records, each record being provided by a data source and having a plurality of terms relating to the object, each data source having an accuracy indicating the accuracy of information provided by the data source;

for each record, calculating a likelihood that the object is relevant to the target term by generating a language model probability for the record; and

calculating the relevance of the object to the target term by combining the likelihoods of the records based on the accuracies of the records.

22. The computer-readable storage medium of claim 21 wherein the accuracy represents accuracy of records of the data source.

23. The computer-readable storage medium of claim 21 wherein the accuracy represents accuracy of attributes of the records of the data source.

24. The computer-readable storage medium of claim 21 wherein each record is represented as a separate document comprising terms of the record.

25. The computer-readable storage medium of claim 21 wherein the language model probability is generated from attributes of each record.

26. The computer-readable storage medium of claim 25 wherein each attribute of a record is represented as separate document comprising terms of the attribute.

27. The computer-readable storage medium of claim 25 wherein the generated language model probability is based on accuracy of records and accuracy of attributes of a data source.

28. A computing system for determining relevance of an object to a target term, comprising:

a memory storing computer-executable instructions of:

a component that retrieves from data sources descriptions of the object, each description including terms used to describe the object;

a component that generates a collection of records of terms relating to the object, the collection including a record generated from each of the plurality of retrieved descriptions, each record including a plurality of terms, each record of the collection being generated by extracting from a retrieved description terms used to describe the object;

a component that, for each record of the collection of records of terms relating to the object, determines a language model probability for that record generating the target term, the language model probability for that record generating the target term being based on a number of occurrences of the target term in that record relative to a number of occurrences of the target term in the collection of record; and

a component that calculates the relevance of the object to the target term based on the language model probability of each of the plurality of the records; and

a processor that executes the computer-executable instructions stored in the memory.

29. The computing system of claim 28 wherein each record is from a data source having an accuracy and wherein the relevance factors in the accuracy of the data sources of the record.

30. The computing system of claim 29 wherein the accuracy represents accuracy of records of the data source.

31. The computing system of claim 29 wherein the accuracy represents accuracy of attributes of the records of the data source.

32. The computing system of claim 28 wherein the component that determines a language model probability factors in probability from attributes of each record.

33. The computing system of claim 32 wherein each attribute of a record is represented as a separate document comprising terms of the attribute.

34. The computing system of claim 32 wherein the determined language model probability is based on accuracy of records and accuracy of attributes of a data source.

35. The computing system of claim 32 wherein the determined language model probability is based on importance of each attribute.

36. The computing system of claim 32 wherein a contribution of each attribute to the determined language model probability for a record factors in importance of each attribute based on the accuracy of the attributes.

37. A computer-readable storage medium storing computer-executable instructions for controlling a computing device to determine relevance of an object to a target term, by a method comprising:

retrieving from data sources descriptions of the object, each description including terms used to describe the object;

generating a collection of records of terms relating to the object, the collection including a record generated from each of the plurality of retrieved descriptions, each record including a plurality of terms, each record of the collection being generated by extracting from a retrieved description terms used to describe the object;

for each record of the collection of records of terms relating to the object, determining a language model probability for that record generating the target term; and calculating the relevance of the object to the target term based on the language model probability of each of the plurality of the records.

38. The computer-readable storage device of claim **37** wherein the calculating of the relevance includes summing the language model probability of each of the plurality of records.

39. The computer-readable storage device of claim **38** wherein the summing weights each language model probability by an accuracy of the data source from which a description is retrieved.

40. The computer-readable storage device of claim **37** wherein the relevance factors in an accuracy of the data source from which a description is retrieved.

* * * * *