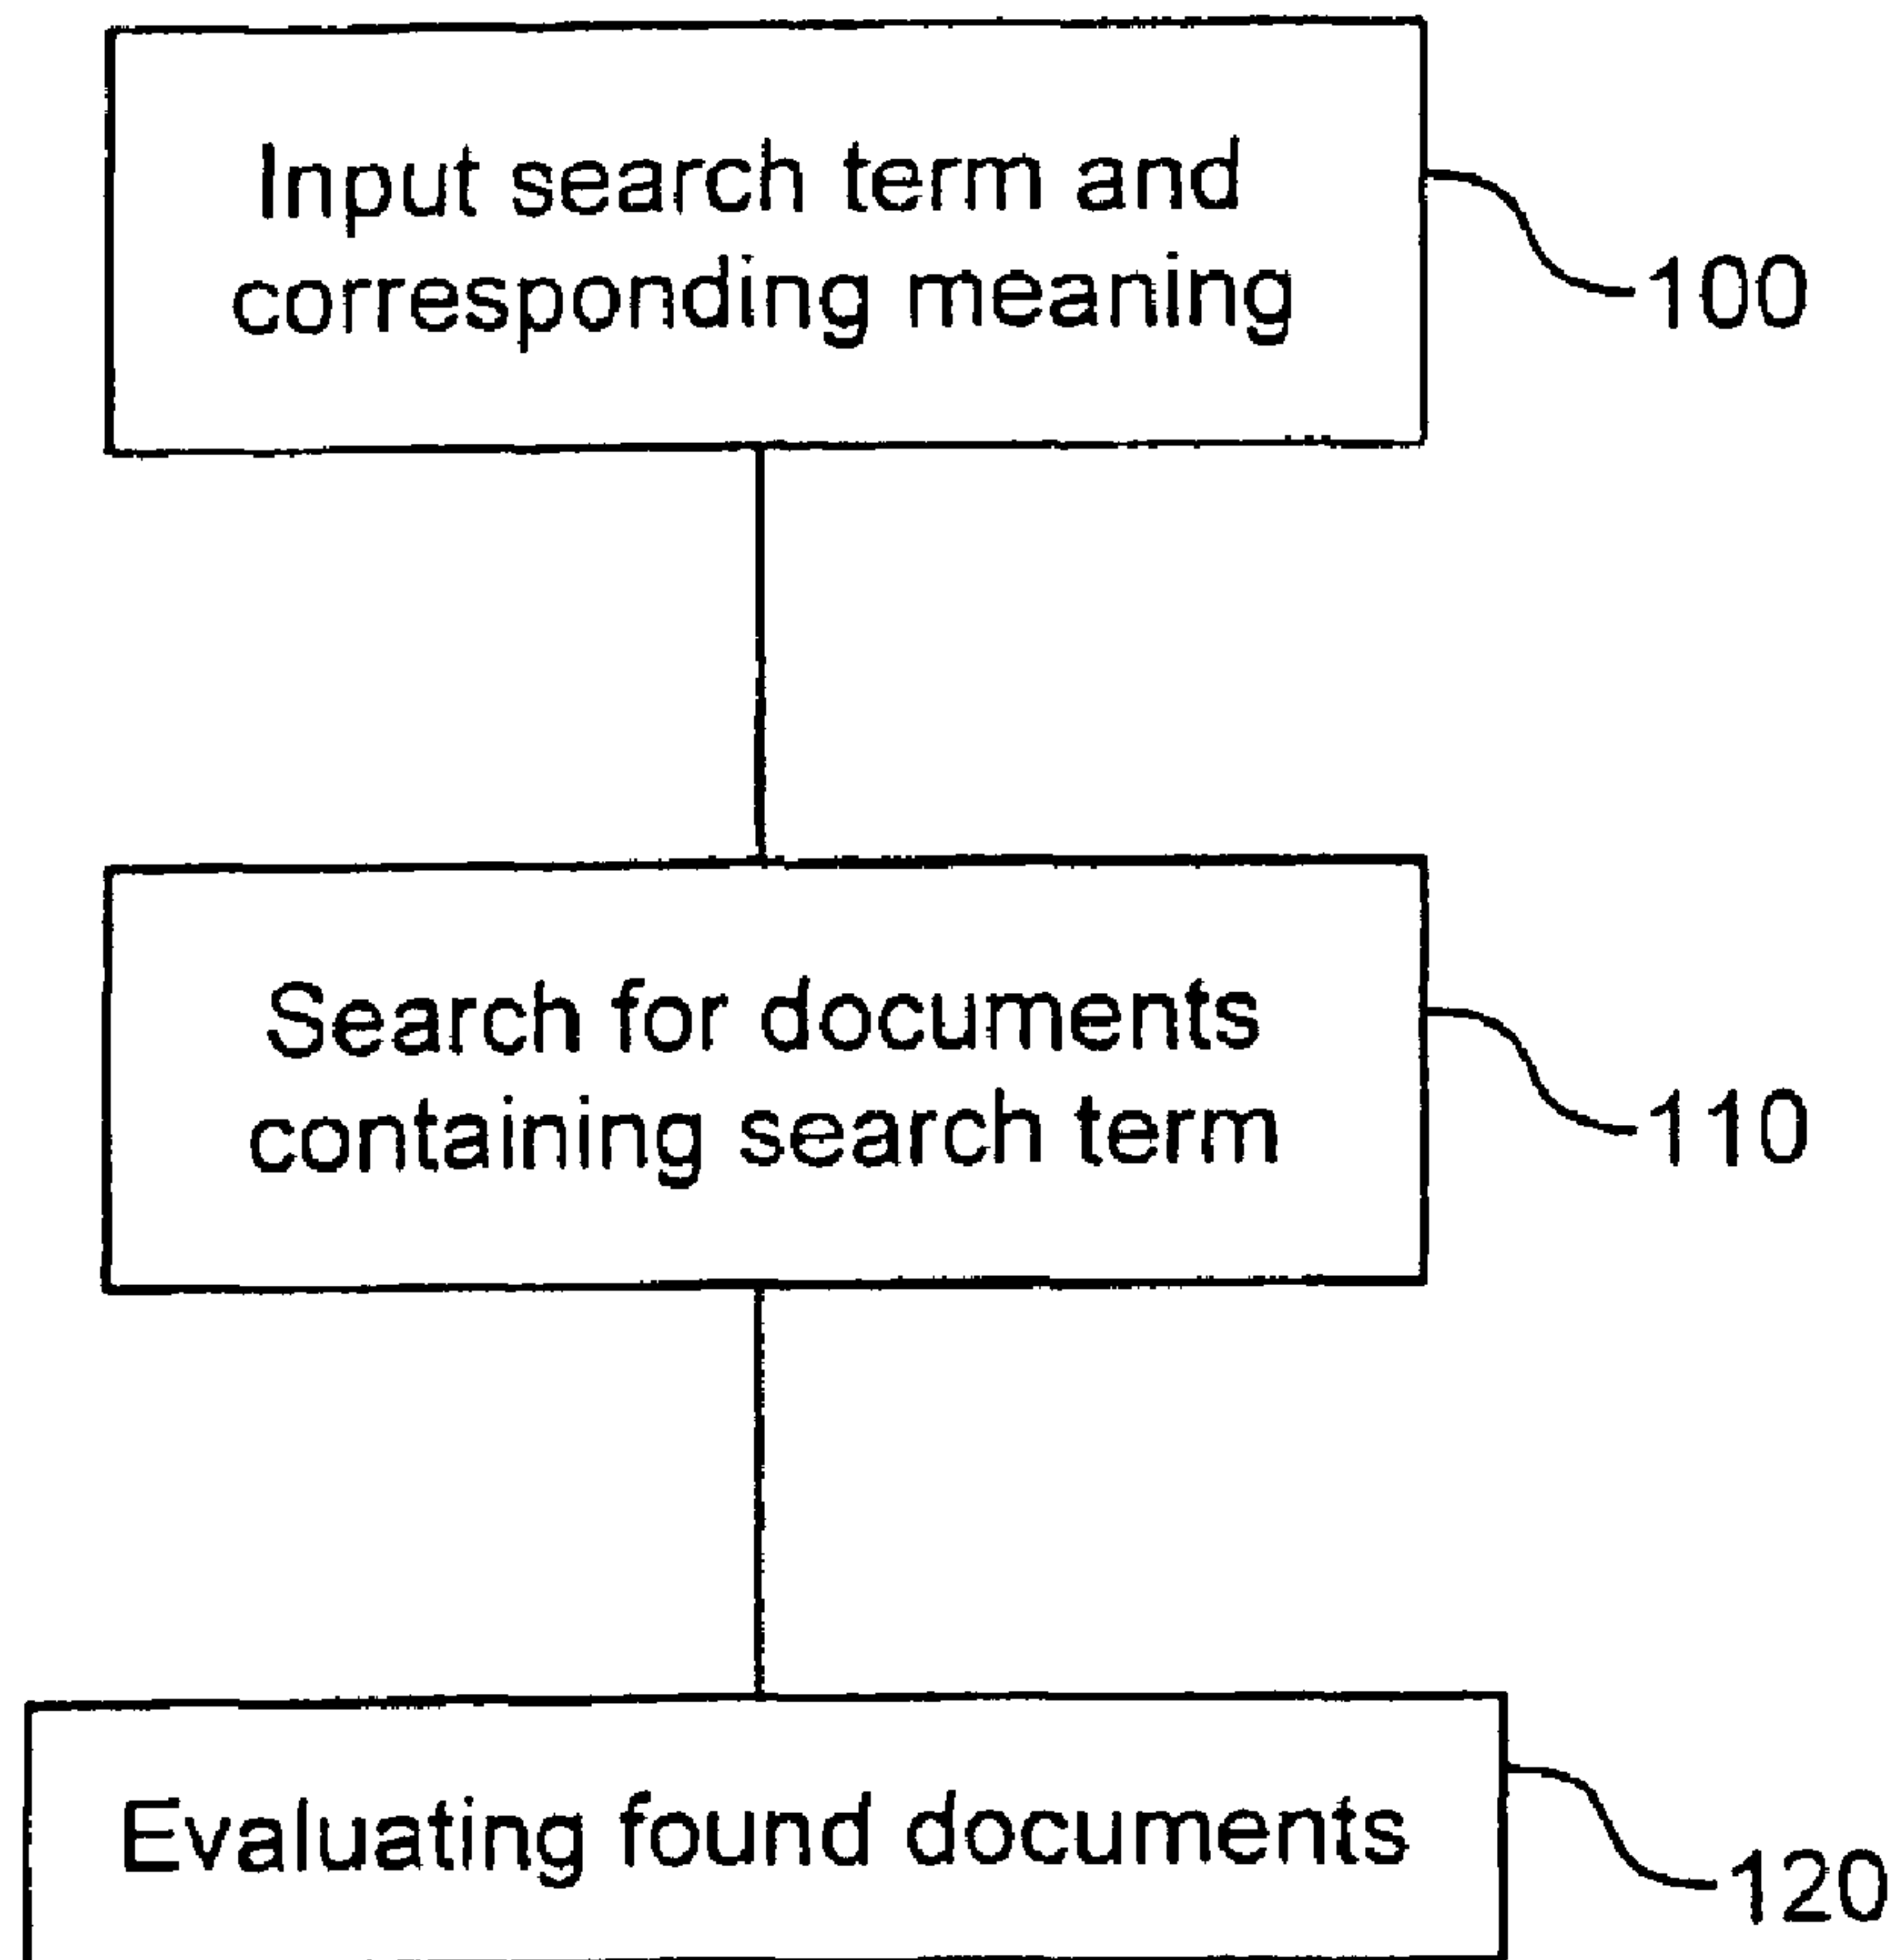




(22) **Date de dépôt/Filing Date:** 2002/08/27  
 (41) **Mise à la disp. pub./Open to Public Insp.:** 2003/03/06  
 (45) **Date de délivrance/Issue Date:** 2015/11/17  
 (62) **Demande originale/Original Application:** 2 459 182  
 (30) **Priorité/Priority:** 2001/08/27 (EP01120429.4)

(51) **Cl.Int./Int.Cl. G06F 17/27** (2006.01),  
**G06F 17/30** (2006.01)  
 (72) **Inventeurs/Inventors:**  
RABALD, MATTHIAS, DE;  
GOERKE, ALEXANDER, DE  
 (73) **Propriétaire/Owner:**  
BDGB ENTERPRISE SOFTWARE SARL, CH  
 (74) **Agent:** NORTON ROSE FULBRIGHT CANADA  
LLP/S.E.N.C.R.L., S.R.L.

(54) **Titre : PROCÉDE D'EVALUATION AUTOMATIQUE DE DOCUMENTS**  
 (54) **Title: A METHOD FOR AUTOMATICALLY EVALUATING DOCUMENTS**



(57) **Abrégé/Abstract:**

A method for evaluating text documents which comprises searching for base documents among a set of base documents which contain a certain search term is described. The method further comprises evaluating the found base documents as to whether certain search term contained in the found base documents, respectively, has a certain meaning.



### Abstract

A method for evaluating text documents which comprises searching for base documents among a set of base documents which contain a certain search term is described. The method further comprises evaluating the found base documents as to whether certain search term contained in the found base documents, respectively, has a certain meaning.

## A METHOD FOR AUTOMATICALLY EVALUATING DOCUMENTS

Cross-reference to related application

This application is a divisional of Application Ser. No. 2,459,182, filed on August 27, 2002.

5 Field of the invention

The present invention relates to the processing of text documents, and in particular, to the automatic evaluating of text documents.

Background of the invention

10 A typical way of structuring large volumes of data such that they can be easily accessed is to index the documents. This means that a document or a group of documents is referenced by an indexing term. A collection of such indexing terms then forms an index. This is exemplarily shown in Fig. 4.

In Fig. 4 address documents 410, 420, 430 contain address data on individual persons. The documents may be characterized by the individual elements  
15 which they contain, one of them is the family name of the persons to which the documents relate.

This information may then be used for building an index 440 shown on the left-hand side of Fig. 4. This index contains list of the family names which are contained in the documents, and each of the elements of the index references an  
20 individual document as can be seen from Fig. 4.

This is a very classical and typical way of organizing information in a structured manner such that the desired information or the desired documents containing a searched and desired piece of information may be retrieved and accessed from a large volume of documents.

25 Indexes can be built for several elements such as the family name, the first name, the street name, etc. What indexes have in common is that the elements

of an index all in some sense have the same "meaning", such as "family name", "first name", or the like.

5 Therefore, the individual elements which are used to build an index are somehow consistent with respect to the information they contain when viewed from a more abstract level. In other words, all elements of the index have the same "meaning".

10 Another, more general approach for ordering documents is just to characterize one or more documents by a certain term, and then to build an index from the thus used individual terms. In such a case the index elements do not have to have a consistent "meaning", although one may consider that they all have the same meaning in the sense that each of the elements characterizes or describes the one or more documents which it references.

15

After an index has been built, it can be used for querying and accessing the set of documents ordered or structured by the thus built up index. Either one can directly input a search term, and if it is contained in the index, then the document or the documents referenced by the index term are retrieved. Another  
20 possibility is to "browse" the index, which means to display the individual index elements in some (typically alphabetical) order, as shown in element 440 of Fig. 4. This has the advantage that a user can quickly overlook which index elements are used in total for organizing or "indexing" the set of documents.

25

Another somewhat more sophisticated approach is to use a so-called fault-tolerant search, which means that a search term is inputted and those documents are retrieved where the corresponding index value is identical or at least similar (to some extent, depending on the used fault-tolerant search algorithm) to the search term.

30

In any case, building an index is a very difficult and tedious work, it is the preparatory work which has to been done in order to make it feasible to access large sets of documents in an ordered and meaningful manner.

5           Typically indexes are created "manually", at least in case of the documents to be indexed being "unstructured", such as plain text documents. If the documents to be indexed are "structured", such as in the case of relational data base tables, then it is relatively easy to built an index. However, if one does not know which individual "meaning" an element in an unstructured document has, then  
10 it is extremely difficult and tiresome to select elements which can be used for indexing this document.

Consequently, it is highly desirable to improve the processing of indexing documents.

15

#### Summary of the invention

According to the first embodiment of present invention there is provided a method for retrieving certain documents from a set of documents based on a  
20 search term inputted as well as based on the corresponding meaning of the search term. At first those documents are searched which contain the search term. Thereafter, the found documents (called base documents) are to be evaluated whether in those documents the search term contained therein has a certain meaning. This evaluation comprises the generation of a text (or layout) document  
25 which represents the search term and its surrounding elements and their corresponding absolute or relative position with respect to the search term or within the base document. This text document then is inputted into a trainable classifying apparatus which has been trained to judge whether the inputted text document belongs to a certain class or category or not. The judgement thereby has to  
30 determine whether the search term has the certain desired meaning (i.e. the text document belongs to the trained class) or whether it has not (the document does not belong to the trained class).

Based on this judgement it can be checked whether in the found based documents the searched terms have the certain desired meaning.

5           According to a further embodiment the search terms and the corresponding meanings can be used to built an index for the based documents. This can then be helpful to access a large stockpile of base documents in an ordered manner.

10           According to a further embodiment searching, classifying, and indexing steps are repeated for a plurality of search terms: If this is carried out for each of the base documents, then for each document the corresponding individual value having a certain meaning and being contained in this document can be determined. For example, for all base documents of a pool of letters the dispatch date may be  
15 determined or "extracted", and the thus extracted elements together with their meaning "dispatch date" could be used to form an index referencing the set of base documents. If this is performed for several different meanings, such as the addressee, the origination address, the reference number, and so on, then there could be also generated a relational data base mapping the "informational content"  
20 of the set of documents into the relational data base tables.

          According to a further embodiment the search performed may comprise an associative search and/or a fuzzy search. This will return also documents in which the search terms are misspelled or which contain any other typing errors.

25

          According to a further embodiment the training of a classifying apparatus can be carried out by selecting for the training set of documents those elements which in the respective documents have a certain meaning, then generate the text documents for the respective base documents of the training sample, and then  
30 using the thus generated text documents as a training set for training the classifying apparatus by running the classifying apparatus in the training mode.

Using this approach a classifying apparatus by using only a small sample of training documents can actually "learn" to evaluate whether for certain base documents some elements contained therein have a certain meaning or not. This so trained "knowledge" can then later on be used for building an index based on the  
5 certain meaning which the index elements should have.

The training process can either be done semi-automatically or it can be done in another embodiment in a fully automatized manner. For the semi-automatic training a user has to identify in the base documents those elements which have  
10 the desired meaning which is to be learned or trained. After the identification has taken place the training sample is provided and the training can be carried out.

To carry out a fully automatic training one can use a set of documents (such as documents of a database which already exists) for which already the  
15 elements having the desired meaning are known, e.g. as fields in the database. This stock of documents for which the elements having the desired meaning are already known can then be used as training samples.

According to a further embodiment a set of base documents can be  
20 automatically indexed by checking for all elements (or at least a part of them) contained in set base documents whether those elements have a certain meaning or not. With this "pump gun" approach a set of base documents can be fully automatically indexed, without even performing any search. In this embodiment just all elements of the base documents are checked as to whether they have a certain  
25 desired meaning, and if so, then those elements are used for building up the index.

According to a further embodiment, the selection of the elements for which the actual checking is performed, may contain some "smart selection". For example, only those elements which meet some format criterion, could be selected  
30 for the evaluation procedure. This could then reduce the necessary computing power to perform the automatic indexing.

According to a further embodiment the one or more indexes which have been built up can be used for forming relational data base tables. Then, the unstructured base documents can be automatically converted into a structured form of a relational data base.

5           According to a further broad aspect of the present invention there is provided a method for evaluating documents. The method comprises searching for base documents among a set of base documents which contain a certain search term. The method further comprises evaluating the found base documents as to whether said certain search term contained in the found base documents, respectively, has a certain meaning.

10           Brief description of the drawings

Fig. 1 schematically illustrates a first embodiment of the present invention.

Fig. 2 schematically illustrates a further embodiment of the present invention.

15           Fig. 3 schematically illustrates an even further embodiment of the present invention.

Fig. 4 schematically illustrates the indexing of documents.

Detailed description

20           The present invention will now be explained by means of illustrative embodiments in connection with the accompanying drawings.

Fig. 1 illustrates a first embodiment according to the present invention. Assume that there is a set of documents which are unstructured (just normal text documents) and which accordingly are difficult to query to access, and to bring them into somehow an ordered form.

25           Let us moreover assume that the user is interested in those documents of the pool of base documents which a) contain a certain term, and b) in which the certain term has a certain meaning. As an example let us assume that a user is interested in those documents of the pool of base documents which contain the

term "May 6, 2000", but additionally, in which the term (or element) "May 6, 2000" is an invoice date.

5 At first a user then has to input the search term into the computer by using any conventional input method. Moreover, the corresponding meaning which the inputted search term should have also must be somehow conveyed to the computer system. Either the user could manually input the corresponding meaning (such as inputting "invoice date", or he could select from a variety of optional meanings, or the corresponding meaning could be predefined to a default value.

10

After the search term and the corresponding meaning have been clarified by operation 100 in Fig. 1, the method proceeds to operation 110. Therein the pool of base documents is searched for those documents which contain the inputted search term.

15

This will retrieve any documents which contain the element "May 6, 2000". However, the retrieved documents may also contain those documents where the search term has a meaning different from the meaning "invoice date", e.g. the meaning could be "date of birth", "date of expiration", or anything else.

20

Therefore, in operation 120 it has to be evaluated whether the found documents are documents in which the search term has the desired corresponding meaning or not.

25 This evaluation procedure will now be explained in more detail in connection with Fig. 2. In operation 200 there is generated a text document with codes for positioning of elements which surround the search term in the found base document. This is explained in very detail in U.S. Patent Application 2008/040660, published on February 14, 2008. In particular, this application deals with and describes a method which can be used for judging whether certain elements in a text have a certain

meaning or not. This is done by representing their corresponding positions through a so-called "layout document", and this layout document then is inputted into a classifying apparatus which has been trained to recognize whether the inputted layout document belongs to a certain category or not.

5

The training usually is performed in order to enable the classifying apparatus to judge whether the layout document represents the surrounding area of a search term for those documents where the search term has any desired (or trained) meaning or not.

10

This layout document (or as we call it here, text document) is then inputted into a classifying apparatus as illustrated in operation 210 of Fig. 2 of the present application. This classifying apparatus then judges whether the inputted text document belongs to a certain category or not. The "category" here means that either the search term has the certain meaning or it has not (here: it is an invoice date or it is not).

15

In operation 220 the classifying apparatus then performs its judgment whether the search term has the inputted (or desired) meaning or not. A particularly suited classifying apparatus is described in European Patent Application 99108354.4, which has been filed by the applicant of the present application on April 28, 1999.

20

In particular, any details relating to the classifying apparatus described or mentioned herein can be taken from this application.

25

However, any other classifying apparatus which is able to evaluate text documents as to whether they belong to a certain category or not and which is trainable can be used as well. Of course, at first the classifying apparatus has to be trained using a certain training sample, and thereafter, the newly found base documents may be inputted into classifying apparatus and evaluated thereby as to whether the search term has a certain meaning or not.

It should be noted once again that a detailed description of the process to extract elements from text documents having a desired meaning is described in European Patent Application 00103810.8 mentioned before. Also the particulars concerning how a representation of the surrounding area of a candidate element  
5 my means of a layout document which is then fed to a trainable classifying apparatus to judge whether the candidate element actually has the desired meaning is described in detail therein.

A further embodiment of the present invention which relates to the  
10 automatic generation of any index is now explained in connection with Fig. 3. In operation 300 the certain meaning which the elements of the base documents should have is defined or selected, as already explained before.

Then, as also already has been explained, for elements to be checked  
15 the text document (or layout document) is generated.

This text documents may then be inputted into an already trained classifying apparatus in operation 320. This classifying apparatus in operation 330 then checks whether the elements to be checked have the certain defined or  
20 selected meaning.

If the evaluation results in a "yes", then in operation 340 the thus evaluated element is used to built an index together with the defined or selected meaning. If the answer in operation 330 is "no", then in operation 350 the checked  
25 elements are discarded or, in other words, they are not used for building the index.

Using the process illustrated in Fig. 3, an automatic indexing of documents can be performed.

30 This can be done in several ways. Either, the elements to be checked could be all elements which are contained in the individual base documents. For each individual element, it could then for example be checked whether this element

is an "invoice date" or not, e.g. by generating a corresponding layout document, and then by evaluating this layout document through a classifying apparatus.

It is readily apparent, that this approach is rather time and computing power consuming, therefore, the selection of the elements to be checked could be done in some smart manner. For example, only those elements which match with a certain format criterion could be checked. Such a format criterion can be a compilation of the well-known formats which a date can assume, and the elements are checked whether they match with the format criterion and only in the latter case the final checking is made whether the thus selected element (which may be assumed to be "date") actually is an "invoice date" or not.

Another approach could be to automatically generate the search terms which then are searched in the base documents.

15

For example, a certain range of dates could be selected (e.g. from January 1, 2000 to January 1, 2001), and then all those dates lying therein between are subsequently searched, in other words, are applied as search terms to the pool of the base documents.

20

For those search terms where there is found a match, which means that there is found a document which contains a corresponding search term, the layout document then is generated and it is checked whether the search term has the certain meaning (here: invoice date). If so, then as already explained before, the search term may be used for building an index.

25

With this approach, not all elements in the base documents have to be checked, but rather only those elements for which the search has delivered a match.

Which approach actually is more efficient in terms of necessary computing power may depend on the individual circumstances.

30

It should be clear that the foregoing examples only are illustrative examples, and that of course the method explained so far also can be applied to any other "meanings" which a certain term or element should have, such as e.g. a name, a street name, a family name, a first name, a city name, or anything else.

5

The most general meaning any element could have with respect to a certain document in which it is contained could be that the element just somehow "describes" the document in which it is contained. For example, the term "invoice" may be contained in an invoice, moreover, it just describes that the corresponding document actually is an invoice. Such elements describing the whole content of a document of course can also be used for building an index.

10

In a further embodiment the training of the classifying apparatus is carried out semi-automatically. E.g. a user identifies for each of a set of documents those elements which have the certain desired meaning such as e.g. the invoice number. Based on this selection inputted by the user then the layout document is generated which can be used as a training input for training the classifying apparatus. In this manner several layout documents representing surrounding areas of elements which are invoice numbers are inputted to the classifying apparatus and thereby the classifying apparatus is trained ("learns") to recognize those elements which are invoice numbers and to distinguish them from other elements which do not have this meaning.

15

20

In a further embodiment such a training procedure can be executed in a fully automated manner. Assume that there is a set of documents (e.g. from any already existing database) which are already stored, evaluated and classified in the form of a relational database which has a column "invoice number" in which for each document the invoice number is already stored. Then there is no need to identify the invoice number manually through a user but rather the information from the database can be used to generate the training input. The database provides the invoice number for each document, this number may then be searched in the document itself and after it has

25

30

been located based on the surrounding area which surrounds the invoice number the layout document is generated. This process may be carried out for each of the documents of the database and then results in a large set of layout documents which can be used as training examples for training the classifying apparatus. Thereby a fully automated  
5 training can be carried out without any need for identifying elements by a user.

It should be clear to any skilled person that this procedure can be carried out not only for "invoice numbers" but for any other elements which have a certain meaning and which are already known. This knowledge about the elements themselves and their  
10 meaning needs not to come from a database, it may also be stored e.g. in a spread sheet table, or it may even be written down on paper in the form of a table which then may be scanned. It does not matter from where the knowledge about the elements themselves and their meaning comes, as soon as it exists it may be used to automatize the training process of the classifying apparatus based on this knowledge.

15 While the present invention has been described hereinbefore by means of exemplary embodiments, it will be readily apparent to the skilled reader that modifications may be carried out without departing from the scope of the invention as defined by the appending claims.

## Claims

1. A method of training a classifying apparatus to retrieve based on a search term together with its corresponding meaning from a set of base documents those documents which contain said search term and in which said search term has said corresponding meaning, said method of training comprising:

looking, utilizing a computer, for base documents in which said search term has said corresponding meaning;

selecting, utilizing the computer, said search term by the user;

repeating said looking and selecting until a sufficient set of base documents has been selected to generate a training sample;

generating, utilizing the computer and the training sample, text documents to represent elements surrounding the search term and the elements' corresponding relative positions with respect to said search term, said elements' relative positions with respect to said search term comprising where the elements are located in the surrounding area of the search term, as compared with where the search term is located;

using, utilizing the computer, said generated text documents as said training set for training said classifying apparatus by running said classifying apparatus in the training mode; and

generating a database from the elements and their corresponding meaning.

2. A method for automatically indexing a set of base documents based on a set of training examples, said automatic indexing comprising:

evaluating, using a computer, said base documents by checking for all elements respectively contained therein which meet predefined criteria, and whether the elements have a corresponding meaning, said evaluation comprising:

for those elements to be checked, generating a text document based on said element to be checked and its surrounding elements coding for the surrounding elements' corresponding relative positions with respect to said element to be checked, said relative positions with respect to said elements to be checked comprising where the

surrounding elements are located as compared with where the elements to be checked are located;

inputting, using the computer, said text documents into a trainable classifying apparatus which has been trained to recognize whether an inputted text document belongs to a corresponding classification category or not, whereas said training has been performed based on a training sample of text documents which have been generated for documents in which the element to be checked surrounded by the surrounding elements has said corresponding meaning; judging, using the computer, by said trainable classifying apparatus whether said element has said corresponding meaning; and

for those base documents where elements have been found to have said corresponding meaning, building, utilizing the computer, an index indexing said large volume of base documents, the index comprising said elements with a corresponding reference to the document in which the elements are contained.

3. A method for automatically indexing a set of base documents based on a set of training examples, said automatic indexing comprising:

evaluating said base documents by checking for some or all elements respectively contained therein whether they have a certain meaning, said evaluation comprising:

for those elements to be checked, generating a text document based on said element to be checked and its surrounding elements coding for their corresponding absolute or relative position with-respect to-said element to be checked;

inputting said text documents into a trainable classifying apparatus which has been trained to recognize whether an inputted text document belongs to a certain classification category or not, whereas said training has been performed based on a training sample of text documents which have been generated for documents in which the element surrounded by the surrounding elements has said certain meaning;

judging by said classifying apparatus whether said element has said certain meaning;

for those base documents where elements have been found to have said certain meaning, using said elements and a corresponding reference to the document in

which they are contained to build an index indexing said large volume of base documents.

4. The method of claim 3, further comprising:

checking all elements of said base documents, or checking only those elements of said base documents which match with one or more predefined criteria.

5. The method of claim 3 or 4, further comprising:

generating a relational database from the index elements and their corresponding meaning.

6. The method as claimed in any one of claims 3 to 5 wherein the method is performed by a classifying apparatus trained to perform the method using a training method comprising:

a) looking for base documents in which an element has a certain meaning;

b) selecting said element by the user; and

repeating steps a) and b) until a sufficient set of base documents has been selected to generate a training sample;

generating the text documents for the respective base documents;

using said generated text documents as a training set for training said classifying apparatus by running said classifying apparatus in the training mode.

7. The method claimed in claim 6 wherein said method further comprises:

using a set of documents for which the elements having the desired predefined meaning is already known to automatically generate text documents as training samples to train said classifying apparatus.

8. A computer program comprising computer program code for enabling a computer to carry out a method according to one of the claims 1 to 7.

9. A system for training a classifying apparatus to retrieve based on a search term together with its corresponding meaning from a set of base documents those documents which contain said search term and in which said search term has said corresponding meaning, said method of training comprising:

a computer configured for:

looking for base documents in which said search term has said corresponding meaning;

selecting said search term by the user;

repeating said looking and selecting until a sufficient set of base documents has been selected to generate a training sample;

generating, utilizing the training sample, text documents to represent elements surrounding the search term and the elements' corresponding relative positions with respect to said search term, said elements' relative positions with respect to said search term comprising where the elements are located in the surrounding area of the search term, as compared with where the search term is located; and

using, utilizing the computer, said generated text documents as said training set for training said classifying apparatus by running said classifying apparatus in the training mode;

generating a database from the elements and their corresponding meaning.

10. A system for automatically indexing a set of base documents based on a set of training examples, said automatic indexing comprising:

a computer configured for:

evaluating said base documents by checking for some or all elements respectively contained therein which meet predefined criteria, and whether the elements have a corresponding meaning, said evaluation comprising:

for those elements to be checked, generating a text document based on said element to be checked and its surrounding elements coding for the surrounding elements' corresponding relative positions with respect to said element to be checked, said relative positions with

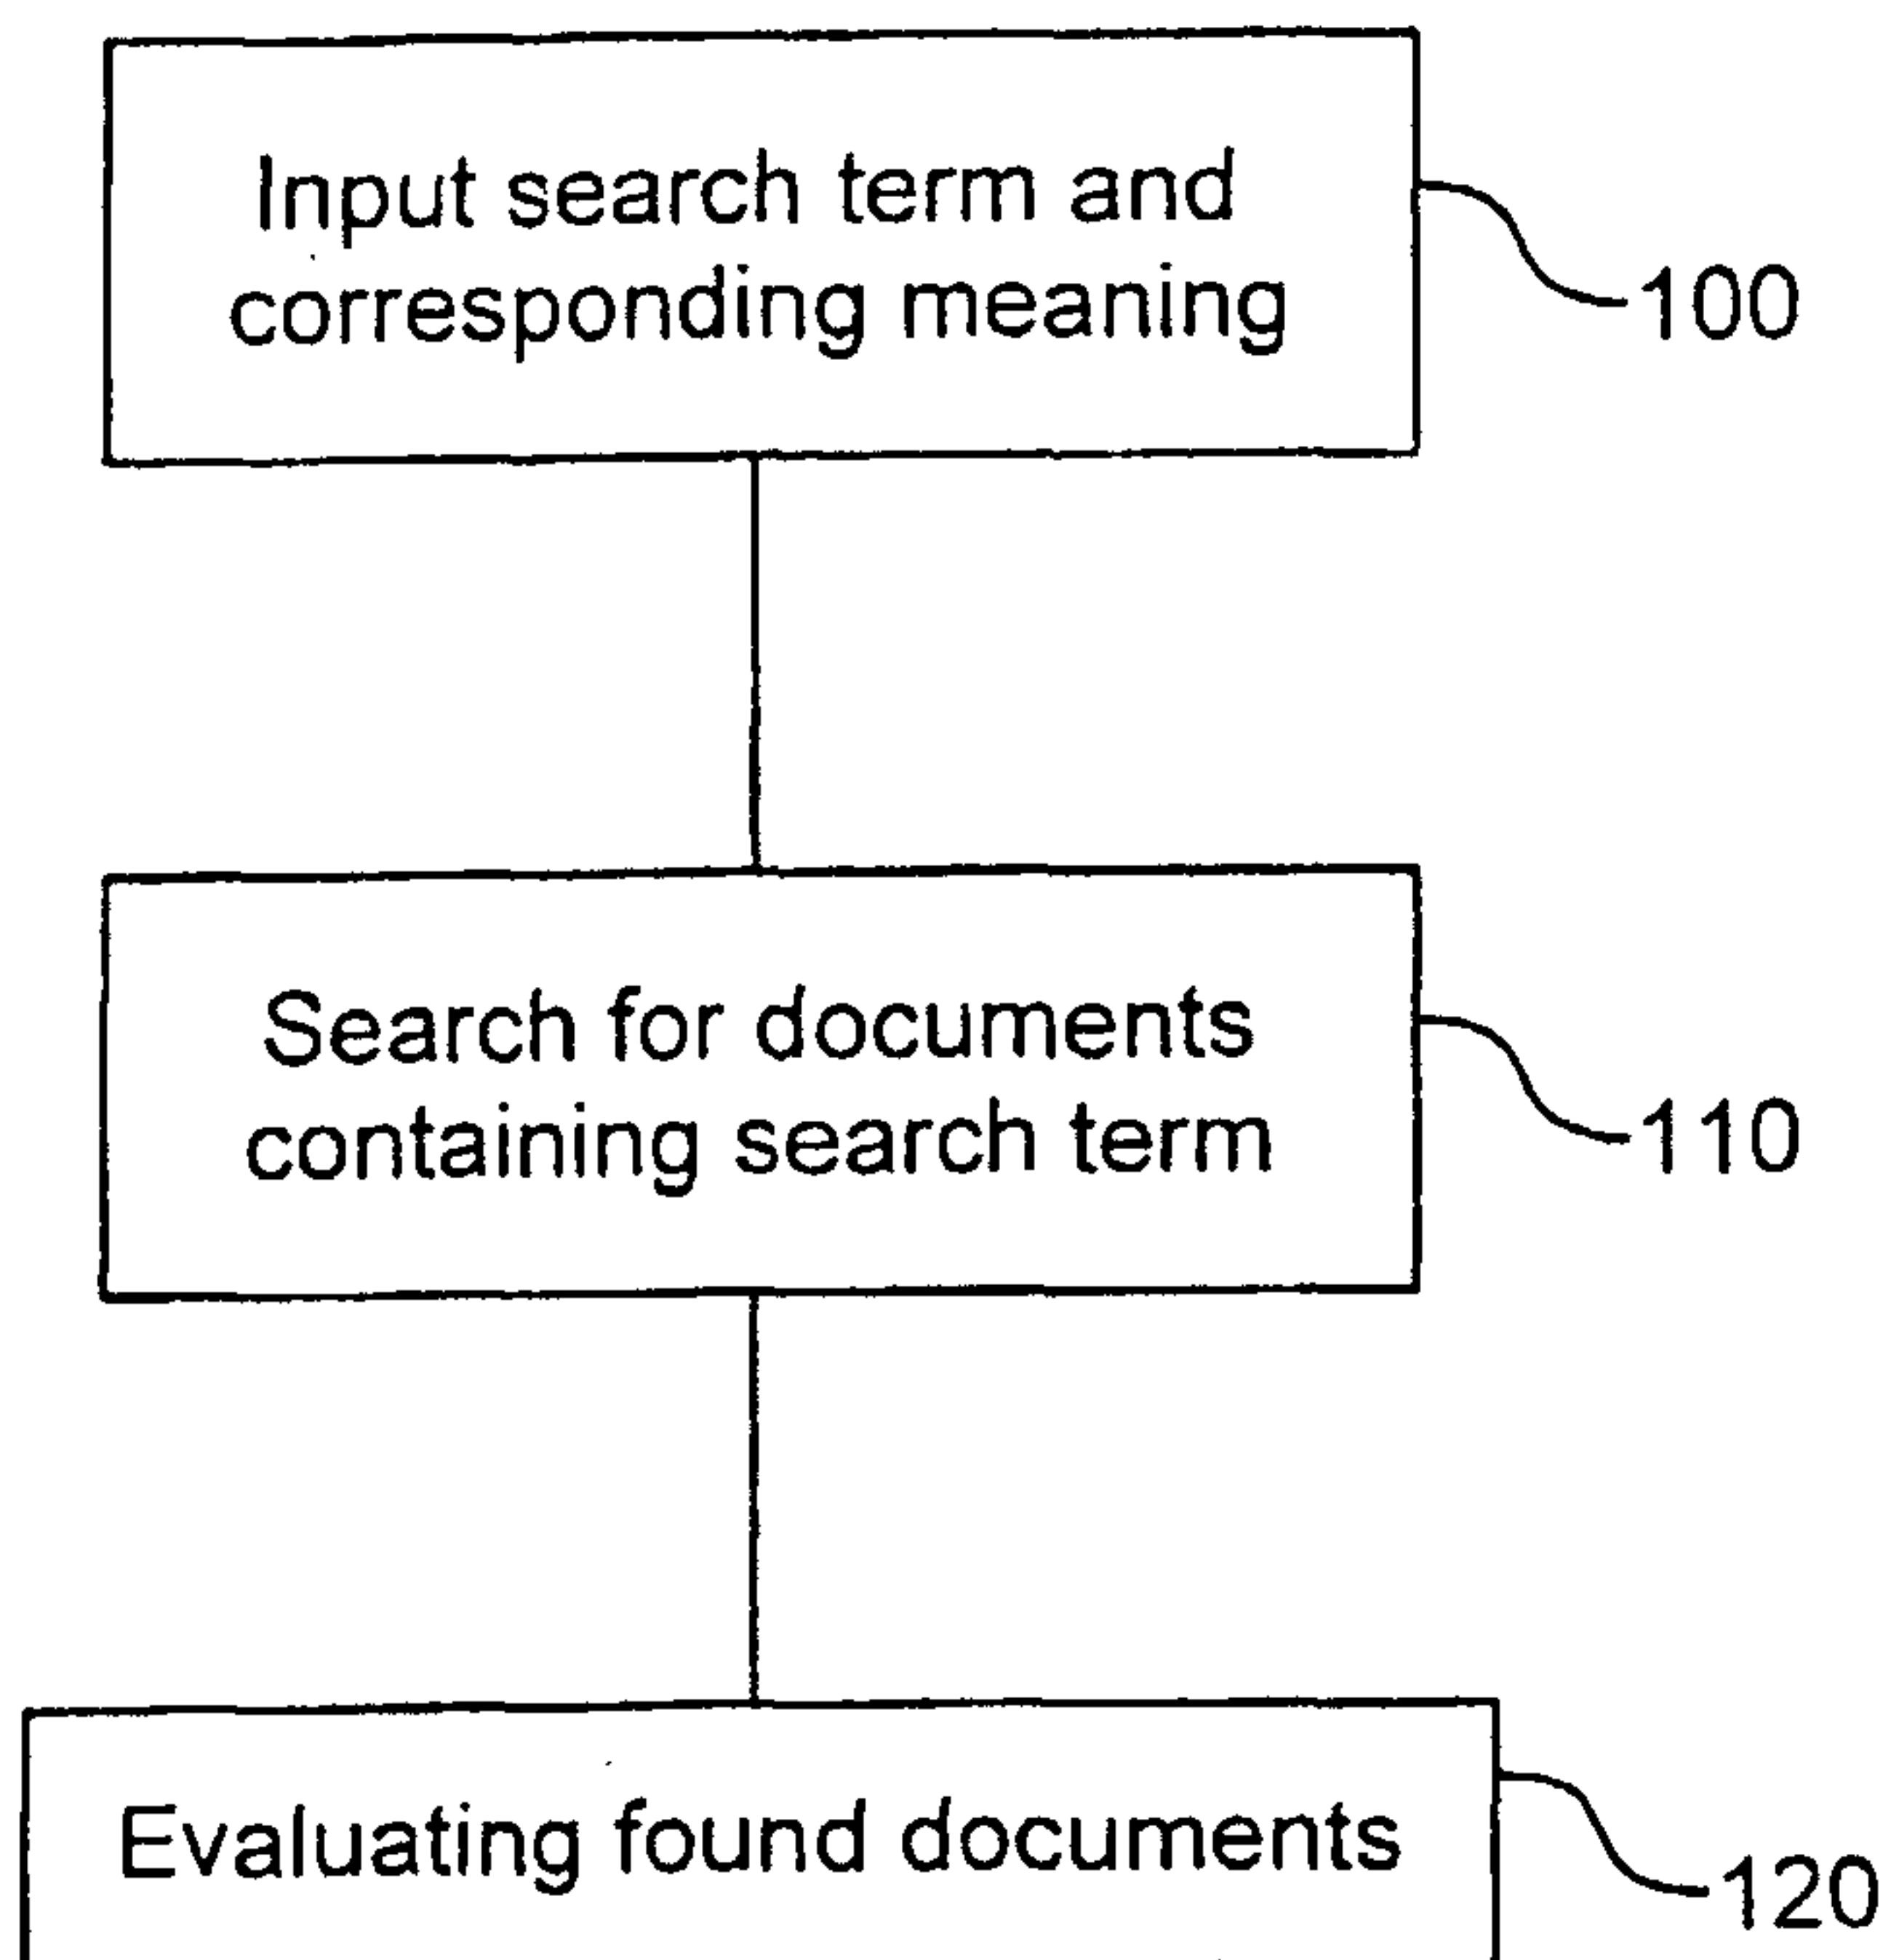
respect to said elements to be checked comprising where the surrounding elements are located as compared with where the elements to be checked are located;

inputting said text documents into a trainable classifying apparatus which has been trained to recognize whether an inputted text document belongs to a predefined classification category or not, whereas said training has been performed based on a training sample of text documents which have been generated for documents in which the element to be checked surrounded by the surrounding elements has said corresponding meaning;

judging by said trainable classifying apparatus whether said element has said corresponding meaning; and

for those base documents where elements have been found to have said corresponding meaning, building, utilizing the computer, an index indexing said large volume of base documents, the index comprising said elements with a corresponding reference to the document in which the elements are contained.

1/4



**Fig. 1**

2/4

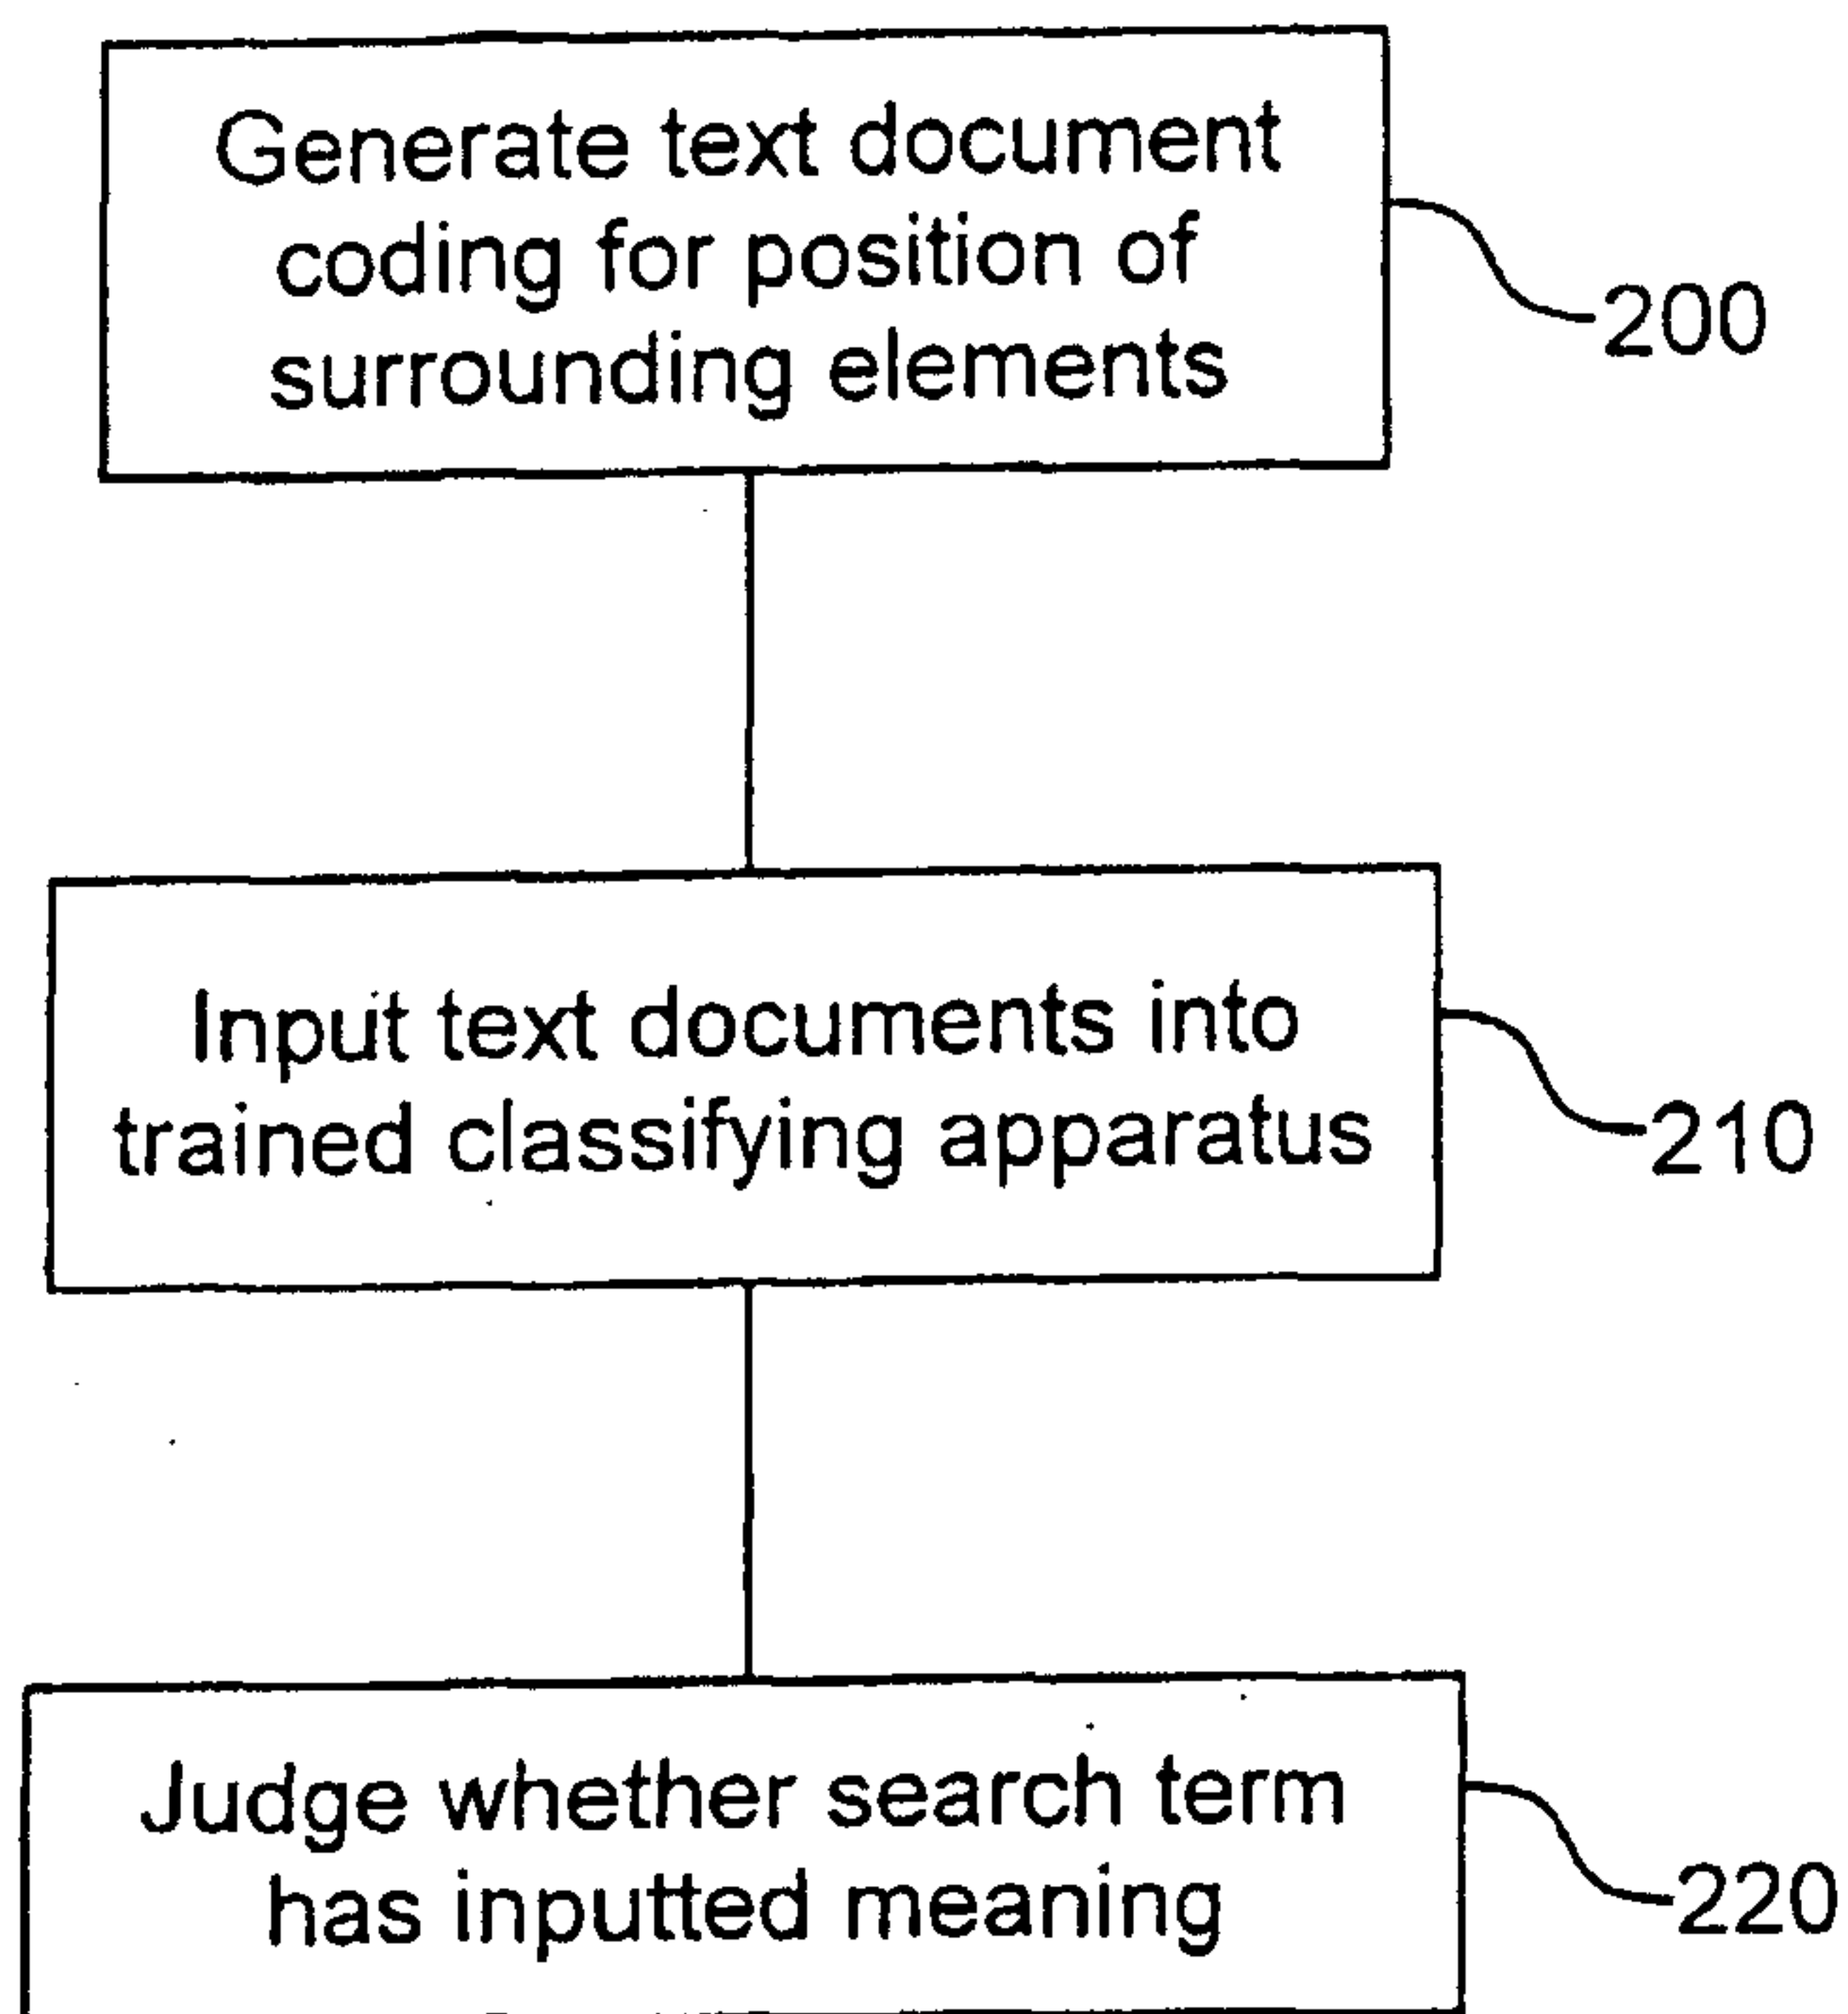
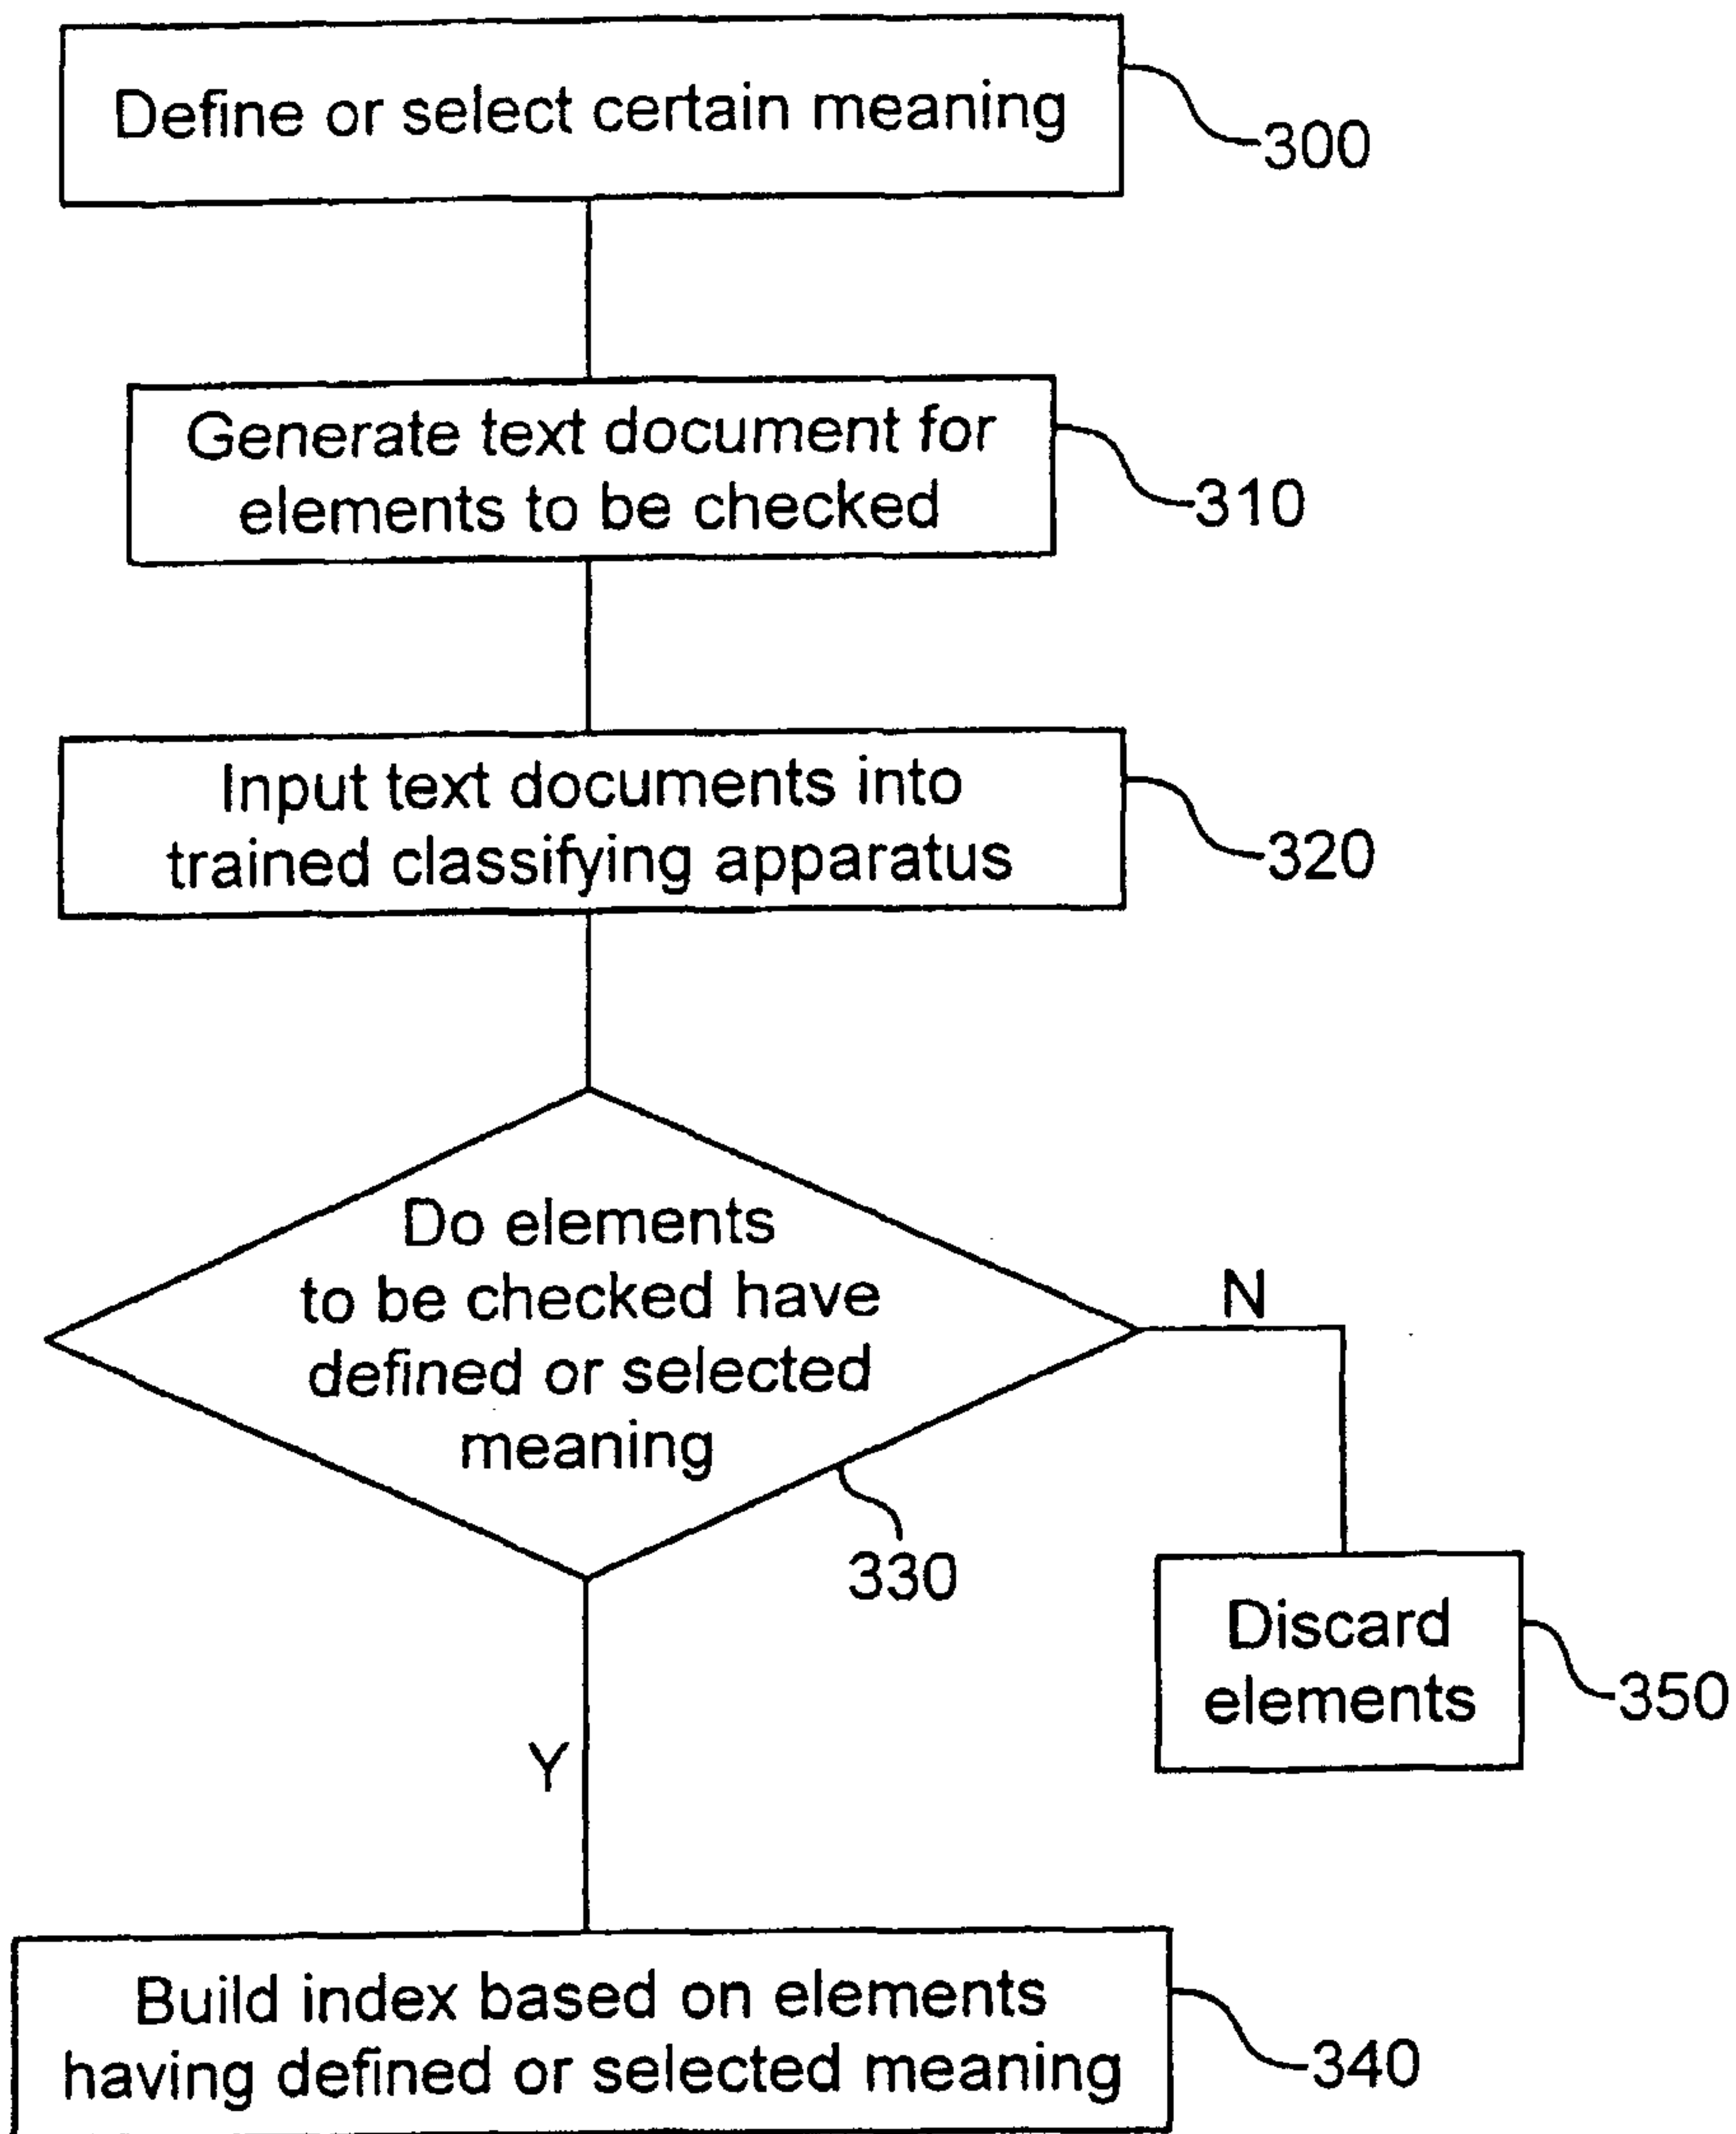


Fig. 2

3/4



**Fig. 3**

4/4

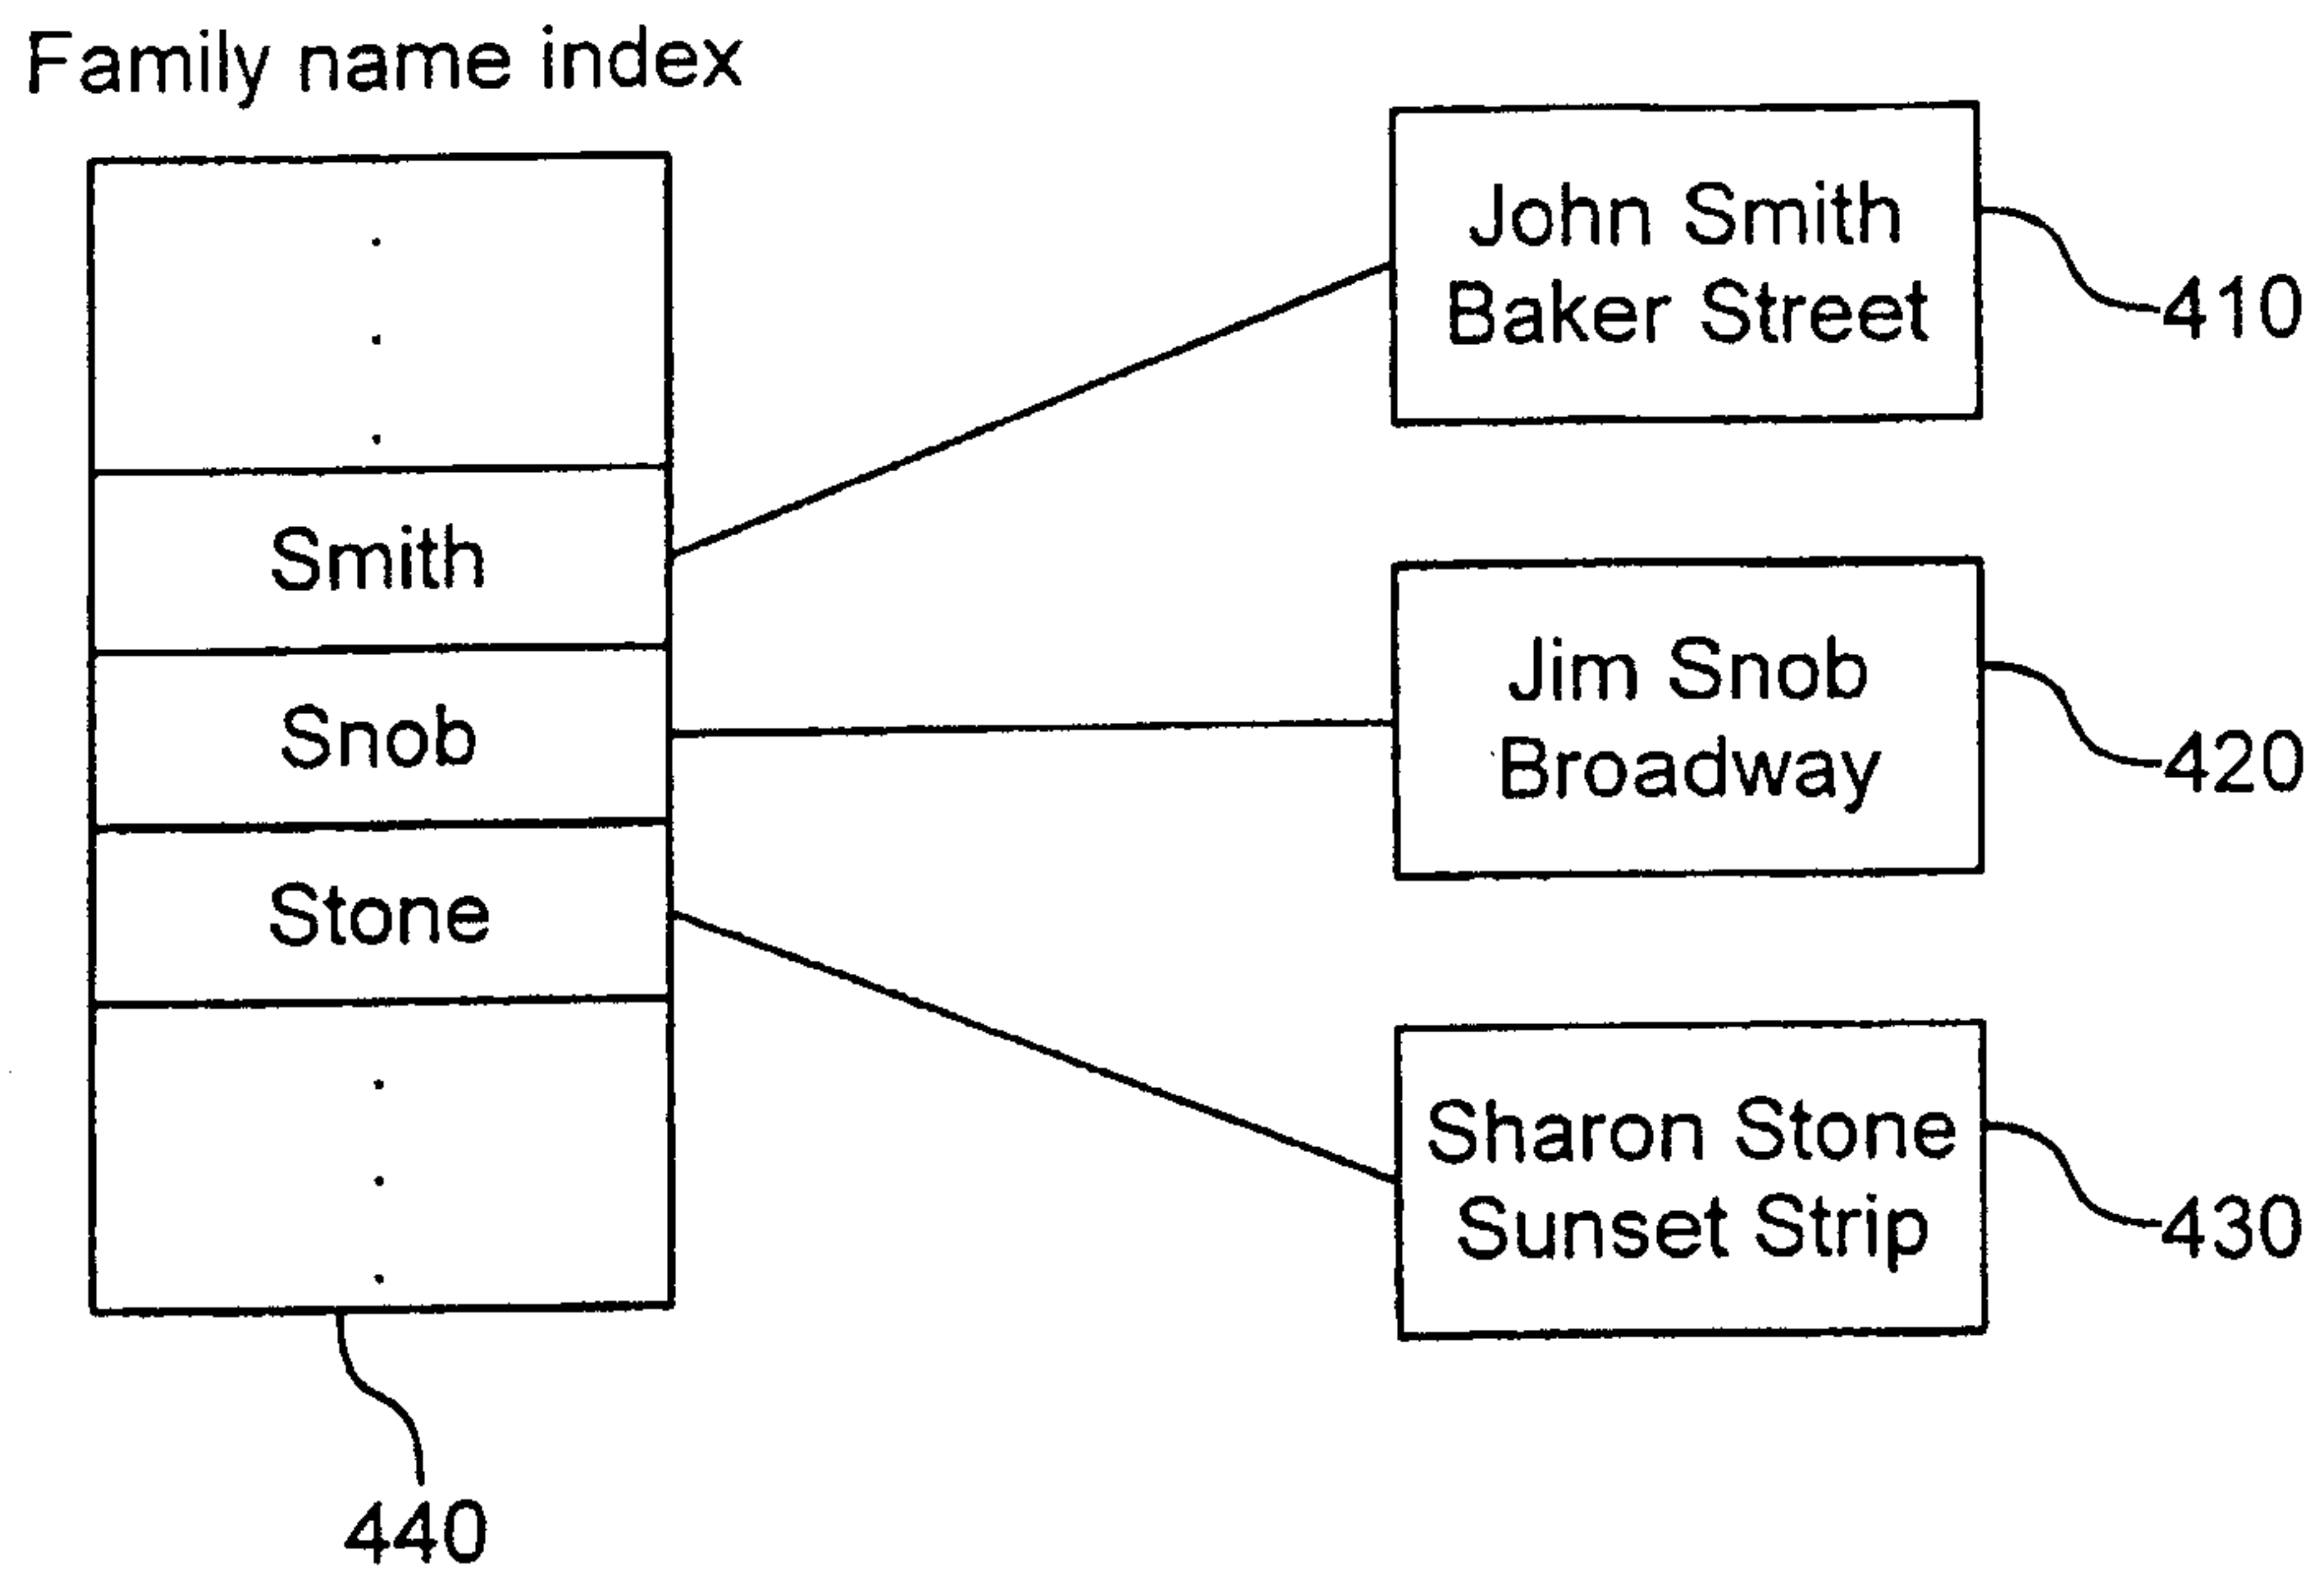


Fig. 4

Input search term and  
corresponding meaning

100

Search for documents  
containing search term

110

Evaluating found documents

120

