

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
21 December 2007 (21.12.2007)

PCT

(10) International Publication Number
WO 2007/146107 A2

(51) International Patent Classification:
G06F 17/30 (2006.01)

(21) International Application Number:

PCT/US2007/013483

(22) International Filing Date: 7 June 2007 (07.06.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/449,400 7 June 2006 (07.06.2006) US

(71) Applicant (for all designated States except US): **DIGITAL MANDATE LLC** [US/US]; 257 Park Avenue South, New York, NY 10010 (US).

(72) Inventor: **KRAFTSOW, Andrew**; 175 Cibola Drive, Sedona, AZ 86336 (US).

(74) Agents: **KUYPER, Reena** et al.; Berry & Associates P.C., 9255 Sunset Blvd., Suite 810, Los Angeles, CA 90069 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, 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).

Declaration under Rule 4.17:

— of inventorship (Rule 4.17(iv))

Published:

— without international search report and to be republished upon receipt of that report

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.

(54) Title: METHODS FOR ENHANCING EFFICIENCY AND COST EFFECTIVENESS OF FIRST PASS REVIEW OF DOCUMENTS

(57) Abstract: Methods for reviewing a collection of documents to identify relevant documents from the collection are provided. A search of the collection can be run based on query terms, to return a subset of responsive documents. A probability of relevancy can be determined for a document in the returned subset, and the document is removed from the subset if it does not reach a threshold probability of relevancy. Documents in a thread of a correspondence (for example, an e-mail) in the responsive documents subset can be added to the responsive documents subset. Further, an attachment to a document in the responsive documents subset can be added to the responsive documents subset. A statistical technique can be applied to determine whether remaining documents in the collection meet a predetermined acceptance level.



WO 2007/146107 A2

**METHODS FOR ENHANCING EFFICIENCY AND COST
EFFECTIVENESS OF FIRST PASS REVIEW OF DOCUMENTS**

FIELD

5 The present disclosure relates to review of documents,
and, more specifically, to techniques for reviewing a
collection of documents to identify relevant documents
from the collection, efficiently and with a relatively
high level of cost effectiveness.

10

BACKGROUND

In the current information age, management of
documents in electronic or paper form can be a
daunting task for an enterprise or other organization.
15 For example in the context of a lawsuit in the United
States, document discovery can entail an enormous task
and large expense, both for the party seeking the
discovery as well as for the party producing documents
in response to document requests from the former.

20

There is a great need for automated methods for
identifying relevant documents. The common method of
discovery today is to round up every document written
or received by named individuals during a time period
25 in question and then read them all to determine
responsiveness to discovery requests. This approach
is obviously prohibitively expensive and time
consuming, and the burden from pursuing such an
approach is increasing in view of the trend of
30 increasing volume of documents.

It has been proposed to use search engine technology
to make the document review process more manageable.
However, the quality and completeness of search

results from conventional search engine techniques are indeterminable and therefore unreliable. For example, one does not know whether the search engine has indeed found every relevant document, at least not with any
5 certainty.

The main search engine technique currently used is keyword or free-text search coupled with indexing of terms in the documents. A user enters a search query
10 consisting of one or a few words or phrases and the search system returns all of the documents that have been indexed as having one or more those words or phrases in the search query. As more documents are indexed, more documents are expected to contain the specified search
15 terms. However, such a search technique only marginally reduces the number of documents to be reviewed, and the large quantities of documents returned cannot be usefully examined by the user. There is no guarantee that the desired information is contained by any of the returned
20 documents.

Further, many of the documents retrieved in a standard search are typically irrelevant because these documents use the searched-for terms in a way or context different
25 from that intended by the user. Words have multiple meanings. One dictionary, for example, lists more than 50 definitions for the word "pitch." We generally do not notice this ambiguity in ordinary usage because the context in which the word appears allows us to pick
30 effortlessly the appropriate meaning of the word for that situation.

In addition, conventional search engine techniques often miss relevant documents because the missed documents do not include the search terms but rather include synonyms of the search terms. That is, the search technique fails to recognize that different words can mean approximately the same thing. For example, "elderly," "aged," "retired," "senior citizens," "old people," "golden-agers," and other terms are used, to refer to the same group of people. A search based on only one of these terms would fail to return a document if the document used a synonym rather than the search term. Some search engines allow the user to use Boolean operators. Users could solve some of the above-mentioned problems by including enough terms in a query to disambiguate its meaning or to include the possible synonyms that might be used.

However, unlike the familiar internet search where one is primarily concerned with finding any document that contains the precise information one is seeking, discovery in a litigation or lawsuit is about finding every document that contains information relevant to the subject. An internet search requires high precision whereas the discovery process requires both high precision and high recall.

For the purposes of discovery in a lawsuit or other legal proceeding, search queries are typically developed with the object of finding every relevant document regardless of the specific nomenclature used in the document. This necessitates developing lists of synonyms and phrases that encompass every imaginable

word usage combination. In practice, the total number of documents returned by these queries is very large.

Methodologies that rely exclusively on technology to
5 determine which documents in a collection are relevant to a lawsuit have not gained wide acceptance regardless of the technology used. These methodologies are often deemed unacceptable because the algorithms used by the machines to determine relevancy are
10 incomprehensible to most parties to a law suit.

There is a need for improved techniques that facilitate the review of a large set of documents, and returns a subset of the documents with a
15 predetermined, high probability that they are relevant.

SUMMARY

This disclosure describes assorted techniques which
20 can be applied in the review of a collection of documents to identify relevant documents from the collection.

A search of the collection can be run based on query
25 terms, to return a subset of responsive documents. In one exemplary embodiment, a probability of relevancy is determined for a document in the returned subset, and the document is removed from the subset if it does not reach a threshold probability of relevancy. A
30 statistical technique can be applied to determine whether remaining documents (that is, not in the

responsive documents subset) in the collection meet a predetermined acceptance level.

5 In another exemplary embodiment, documents in a thread of a correspondence (for example, an e-mail) in the responsive documents subset can be added to the responsive documents subset. the responsive documents in the responsive documents subset are scanned to automatically identify a correspondence (for example,
10 an e-mail) in the responsive documents subset, additional documents in a thread of the correspondence automatically identified, and the additional documents are added to the responsive documents subset.

15 In another exemplary embodiment, the responsive documents in the responsive documents subset are scanned to automatically determine whether any of the responsive documents include an attachment that is not in the subset, and any such attachment is added to the
20 responsive documents subset

In another exemplary embodiment, (a) a predetermined number of documents are randomly selected from a remainder of the collection of documents not in the
25 responsive documents subset, (b) the randomly selected documents are reviewed to determine whether the randomly selected documents include additional relevant documents, (c) if there are additional relevant documents, one or more specific terms in the
30 additional responsive documents that render the documents relevant are identified, the query terms are

expanded with the specific terms, and the search is re-run with the expanded query terms.

BRIEF DESCRIPTION OF THE DRAWINGS

5 The features of the present application can be more readily understood from the following detailed description with reference to the accompanying drawings wherein:

10 **Fig. 1:** A block diagram of a computer or information terminal on which programs can run to implement the methodologies of this disclosure.

Fig. 2: A flow chart for a method for reviewing a
15 collection of documents to identify relevant documents from the collection, according to an exemplary embodiment.

Fig. 3: A flow chart for a method for reviewing a
20 collection of documents to identify relevant documents from the collection, according to another exemplary embodiment.

Fig. 4: A flow chart for a method for reviewing a
25 collection of documents to identify relevant documents from the collection, according to another exemplary embodiment.

Fig. 5: A flow chart for a method for reviewing a
30 collection of documents to identify relevant documents from the collection, according to another exemplary embodiment.

Fig. 6: A flow chart for a method for reviewing a collection of documents to identify relevant documents from the collection, according to another exemplary embodiment.

Figs. 7A and 7B: A flow chart for a workflow of a process including application of some of the techniques discussed herein.

10

DETAILED DESCRIPTION

Non-limiting details of exemplary embodiments are described below, including discussions of theory and experimental simulations which are set forth to aid in an understanding of this disclosure but are not intended to, and should not be construed to, limit in any way the claims which follow thereafter.

20 Full citations for a number of publications may be found immediately preceding the claims. The disclosures of these publications are hereby incorporated by reference into this application in order to more fully describe the state of the art as
25 of the date of the methods and apparatuses described and claimed herein. In order to facilitate an understanding of the discussion which follows one may refer to the publications for certain frequently occurring terms which are used herein.

30

One or more computer programs may be included in the implementation of the apparatuses and methodologies of

this application. The computer programs may be stored in a machine-readable program storage device or medium and/or transmitted via a computer network or other transmission medium. An exemplary configuration of a computer on which the programs can run is shown in Fig. 1. Computer 10 includes CPU 11, program and data storage 12, hard disk (and controller) 13, removable media drive (and controller) 14, network communications controller 15 (for communications through a wired or wireless network), display (and controller) 16 and I/O controller 17, all of which are connected through system bus 19.

Some examples of methodologies, in accordance with this disclosure, for reviewing a collection of documents to identify relevant documents from the collection will now be discussed.

In one example (Fig. 2), a method for reviewing a collection of documents to identify relevant documents from the collection can comprise running a search of the collection of documents based on a plurality of query terms and returning a subset of responsive documents from the collection (step S21), determining a corresponding probability of relevancy for each document in the responsive documents subset (step S23) and removing from the responsive documents subset, documents that do not reach a threshold probability of relevancy (step S25).

Some additional features which are optional include the following.

The techniques discussed in this disclosure are preferably automated as much as possible.

5 Therefore, the search is preferably applied through a search engine. The search can include a concept search, and the concept search is applied through a concept search engine. Such searches and other automated steps or actions can be coordinated through
10 appropriate programming, as would be appreciated by one skilled in the art.

The probability of relevancy of a document can be scaled according to a measure of obscurity of the
15 search terms found in the document.

The method can further comprise randomly selecting a predetermined number of documents from a remaining subset of the collection of documents not in the
20 responsive documents subset, and determining whether the randomly selected documents include additional relevant documents, and in addition, optionally, identifying one or more specific terms in the additional relevant documents that render the
25 documents relevant, expanding the query terms with the specific terms, and re-running at least the search with the expanded query terms. If the randomly selected documents include one or more additional relevant documents, the query terms can be expanded
30 and the search re-run with the expanded query terms. The method can additionally comprise comparing a ratio of the additional relevant documents and the randomly

selected documents to a predetermined acceptance level, to determine whether to apply a refined set of query terms.

- 5 The method can further comprise selecting two or more search terms, identifying synonyms of the search terms, and forming the query terms based on the search terms and synonyms.
- 10 The method can further comprise identifying a correspondence between a sender and a recipient, in the responsive documents subset, automatically determining one or more additional documents which are in a thread of the correspondence, the additional
- 15 documents not being in the responsive documents subset, and adding the additional documents to the responsive documents subset. The term "correspondence" is used herein to refer to a written or electronic communication (for example, letter, memo, e-mail, text
- 20 message, etc.) between a sender and a recipient, and optionally with copies going to one or more copy recipients.

The method can further comprise determining whether

25 any of the documents in the responsive documents subset includes an attachment that is not in the responsive documents subset, and adding the attachment to the responsive documents subset.

- 30 The method can further comprise applying a statistical technique (for example, zero-defect testing) to determine whether remaining documents not in the

responsive documents set meets a predetermined acceptance level.

The search can include (a) a Boolean search of the
5 collection of documents based on the plurality of
query terms, the Boolean search returning a first
subset of responsive documents from the collection,
and (b) a second search by applying a recall query
based on the plurality of query terms to remaining
10 ones of the collection of documents which were not
returned by the Boolean search, the second search
returning a second subset of responsive documents in
the collection, and wherein the responsive documents
subset is constituted by the first and second subsets.
15 The first Boolean search may apply a measurable
precision query based on the plurality of query terms.
The method can optionally further include
automatically tagging each document in the first
subset with a precision tag, reviewing the document
20 bearing the precision tag to determine whether the
document is properly tagged with the precision tag,
and determining whether to narrow the precision query
and rerun the Boolean search with the narrowed query
terms. The method can optionally further comprise
25 automatically tagging each document in the second
subset with a recall tag, reviewing the document
bearing the recall tag to determine whether the
document is properly tagged with the recall tag, and
determining whether to narrow the recall query and
30 rerun the second search with the narrowed query terms.
The method can optionally further include reviewing
the first and second subsets to determine whether to

modify the query terms and rerun the Boolean search and second search with modified query terms.

In another example (Fig. 3), a method for reviewing a collection of documents to identify relevant documents from the collection includes running a search of the collection of documents, based on a plurality of query terms, the search returning a subset of responsive documents in the collection (step S31), automatically identifying a correspondence between a sender and a recipient, in the responsive documents subset (step S33), automatically determining one or more additional documents which are in a thread of the correspondence, the additional documents not being in the responsive documents subset (step S35), and adding the additional documents to the responsive documents subset (step S37).

Some additional features which are optional include the following.

The method can further comprise determining for each document in the responsive documents subset, a corresponding probability of relevancy, and removing from the responsive documents subset documents that do not reach a threshold probability of relevancy. The probability of relevancy of a document can be scaled according to a measure of obscurity of the search terms found in the document.

The method can further comprise applying a statistical technique to determine whether a remaining subset of

the collection of documents not in the responsive documents subset meets a predetermined acceptance level.

5 The method can additionally comprise randomly selecting a predetermined number of documents from a remainder of the collection of documents not in the responsive documents subset, determining whether the randomly selected documents include additional
10 relevant documents, identifying one or more specific terms in the additional relevant documents that render the documents relevant, expanding the query terms with the specific terms, and re-running the search with the expanded query terms.

15

The method can further include randomly selecting a predetermined number of documents from a remainder of the collection of documents not in the responsive documents subset, determining whether the randomly
20 selected documents include additional relevant documents, comparing a ratio of the additional relevant documents and the randomly selected documents to a predetermined acceptance level, and expanding the query terms and rerunning the search with the expanded
25 query terms, if the ratio does not meet the predetermined acceptance level.

The method can further comprise selecting two or more search terms, identifying synonyms of the search
30 terms, and forming the query terms based on the search terms and synonyms.

The method can additionally include determining whether any of the responsive documents in the responsive documents subset includes an attachment that is not in the subset, and adding the attachment
5 to the subset.

In another example (Fig. 4), a method for reviewing a collection of documents to identify relevant documents from the collection can comprise running a search of
10 the collection of documents, based on a plurality of query terms, the search returning a subset of responsive documents in the collection (step S41), automatically determining whether any of the responsive documents in the responsive documents
15 subset includes an attachment that is not in the subset (step S43), and adding the attachment to the responsive documents subset (step S45).

Some additional features which are optional include
20 the following.

The method can further comprise determining for each document in the responsive documents subset, a corresponding probability of relevancy, and removing
25 from the responsive documents subset documents that do not reach a threshold probability of relevancy. The probability of relevancy of a document is preferably scaled according to a measure of obscurity of the search terms found in the document.

30

The method can additionally comprise applying a statistical technique to determine whether a remaining

subset of the collection of documents not in the responsive documents subset meets a predetermined acceptance level.

5 The method can further include randomly selecting a predetermined number of documents from a remainder of the collection of documents not in the responsive documents subset, determining whether the randomly selected documents include additional relevant
10 documents, identifying one or more specific terms in the additional responsive documents that render the documents relevant, expanding the query terms with the specific terms, and re-running the search with the expanded query terms.

15

The method can further include selecting two or more search terms, identifying synonyms of the search terms, and forming the query terms based on the search terms and synonyms.

20

The method can further comprise identifying a correspondence between a sender and a recipient, in the responsive documents subset, automatically determining one or more additional documents which are
25 in a thread of the correspondence, the additional documents not being in the responsive documents subset, and adding the additional documents to the responsive documents subset.

30 In another example (Fig. 5), a method for reviewing a collection of documents to identify relevant documents from the collection comprises running a search of the

collection of documents, based on a plurality of query terms, the search returning a subset of responsive documents from the collection (step S51), randomly selecting a predetermined number of documents from a remainder of the collection of documents not in the responsive documents subset (step S52), determining whether the randomly selected documents include additional relevant documents (step S53), identifying one or more specific terms in the additional responsive documents that render the documents relevant (step S54), expanding the query terms with the specific terms (step S55), and re-running the search with the expanded query terms (step S56).

15 In another example (Fig. 6), a method for reviewing a collection of documents to identify relevant documents from the collection can comprise specifying a set of tagging rules to extend query results to include attachments and email threads (step S61), expanding search query terms based on synonyms (step S62), running a precision Boolean search of the collection of documents, based on two or more search terms and returning a first subset of potentially relevant documents in the collection (step S63), calculating the probability that the results of each Boolean query are relevant by multiplying the probability of relevancy of each search term, where those individual probabilities are determined using an algorithm constructed from the proportion of relevant synonyms for each search term (step S64), applying a recall query based on the two or more search terms to run a second concept search of remaining ones of the

collection of documents which were not returned by the first Boolean search, the second search returning a second subset of potentially relevant documents in the collection (step S65), calculating the probability
5 that each search result in the recall query is relevant to a given topic based upon an ordering of the concept search results by relevance to the topic by vector analysis (step S66), accumulating all search results that have a relevancy probability of greater
10 than 50% into a subset of the collection (step S67), randomly selecting a predetermined number of documents from the remaining subset of the collection and determining whether the randomly selected documents include additional relevant documents (step S68), if
15 additional relevant documents are found (step S69, yes), identifying the specific language that causes relevancy, and expanding that language into a set of queries (step S70), constructing and running precision Boolean queries of the entire document collection
20 above (step S71).

The following discussions of theory and exemplary embodiments are set forth to aid in an understanding of the subject matter of this disclosure but are not
25 intended to, and should not be construed to, limit in any way the invention as set forth in the claims which follow thereafter.

As discussed above, one of the problems with using
30 conventional search engine techniques in culling a collection of documents is that such techniques do not meet the requirements of recall and precision.

However, by using statistical sampling techniques it is possible to state with a defined degree of confidence the percentage of relevant documents that may have been missed. Assuming the percentage missed is set low enough (1%) and the confidence level is set high enough (99%), this statistical approach to identifying relevant documents would likely satisfy most judges in most jurisdictions. The problem then becomes how to select a subset of the document collection that is likely to contain all responsive documents. Failure to select accurately results in an endless cycle of statistical testing.

The probability that results of a simple Boolean search (word search) are relevant to a given topic is directly related to the probability that the query terms themselves are relevant, i.e. that those terms are used within a relevant definition or context in the documents. Similarly, the likelihood that a complex Boolean query will return relevant documents is a function of the probability that the query terms themselves are relevant.

For example, the documents collected for review in today's lawsuits contain an enormous amount of email. It has been found that corporate email is not at all restricted to "business as such" usage. In fact, it is hard to distinguish between personal and business email accounts based on subject matter. As a consequence, even though a particular word may have a particular meaning within an industry, the occurrence

of that word in an email found on a company server does not guarantee that it has been used in association with its "business" definition.

5 An exemplary method for determining a probability of relevancy to a defined context is discussed below.

The following factors can be used to determine the probability that a word has been used in the defined
10 context within a document: (1) the number of possible definitions of the word as compared to the number of relevant definitions; and (2) the relative obscurity of relevant definitions as compared to other definitions.

15

Calculation of the first factor is straightforward. If a word has five potential definitions (as determined by a credible dictionary) and if one of those definitions is responsive, then the basic probability
20 that word is used responsively in any document retrieved during discovery is 20% (1/5). This calculation assumes, however, that all definitions are equally common, that they are all equally likely to be chosen by a writer describing the subject matter. Of
25 course, that is generally not the case; some definitions are more "obscure" than others meaning that users are less likely to chose the word to impart that meaning. Thus, a measure of obscurity must be factored into the probability calculation.

30

A social networking approach can be taken to measure obscurity. The following method is consistent with the

procedure generally used in the legal field currently for constructing query lists: (i) a list of potential query terms (keywords) is developed by the attorney team; (ii) for each word, a corresponding list of synonyms is created using a thesaurus; (iii) social network is drawn (using software) between all synonyms and keywords; (iv) a count of the number of ties at each node in the network is taken (each word is a node); (v) an obscurity factor is determined as the ratio between the number of ties at any word node and the greatest number of ties at any word node, or alternatively their respective z scores; and (vi) this obscurity factor is applied to the definitional probability calculated above.

15

The method described above calculates the probability that a given word is used in a relevant manner in a document. Boolean queries usually consist of multiple words, and thus a method of calculating the query terms interacting with each other is required.

20

The simplest complex queries consist of query terms separated by the Boolean operators AND and/or OR. For queries separated by an AND operator, the individual probabilities of each word in the query are multiplied together to yield the probability that the complex query will return responsive results. For query terms separated by an OR operator, the probability of the query yielding relevant results is equal to the probability of the lowest ranked search term in the query string.

30

Query words strung together within quotation marks are typically treated as a single phrase in Boolean engines (i.e. they are treated as if the string is one word). A document is returned as a result if and only
5 if the entire phrase exists within the document. For purposes of calculating probability, the phrase is translated to its closest synonym and the probability of that word is assigned to the phrase. Moreover, since a phrase generally has a defined part of speech
10 (noun, verb, adjective, etc.), when calculating probability one considers only the total number of possible definitions for that part of speech, thereby reducing the denominator of the equation and increasing the probability of a responsive result.

15

Complex Boolean queries can take the form of "A within X words B", where A and B are query terms and X is the number of words in separating them in a document which is usually a small number. The purpose of this type of
20 query, called a proximity query, is to define the terms in relation to one another. This increases the probability that the words will be used responsively. The probability that a proximity query will return responsive documents equals the probability of the
25 highest query term in the query will be responsive.

A workflow of a process including application of some of the techniques discussed herein, according to one example, is shown exemplarily in Figs. 7A and 7B.

30 The specific embodiments and examples described herein are illustrative, and many variations can be

introduced on these embodiments and examples without departing from the spirit of the disclosure or from the scope of the appended claims. For example, features of different illustrative embodiments and examples may be combined with each other and/or substituted for each other within the scope of this disclosure and appended claims.

References

10 Herbert L. Roitblat, "Electronic Data Are Increasingly Important To Successful Litigation" (November 2004).

Herbert L. Roitblat, "Document Retrieval" (2005).

15 "The Sedona Principles: Best Practices Recommendations & Principles for Addressing Electronic Document Production" (July 2005 Version).

WHAT IS CLAIMED IS:

1. A method for reviewing a collection of documents to identify relevant documents from the collection, the method comprising:

5 running a search of the collection of documents, the search being based on a plurality of query terms and returning a subset of responsive documents from the collection;

determining a corresponding probability of
10 relevancy for each document in the responsive documents subset; and

removing from the responsive documents subset, documents that do not reach a threshold probability of relevancy.

15

2. The method of claim 1, wherein the probability of relevancy of a document is scaled according to a measure of obscurity of the search terms found in the document.

20

3. The method of claim 1, wherein the search is applied through a search engine.

4. The method of claim 1, wherein the search
25 includes a concept search, and the concept search is applied through a concept search engine.

5. The method of claim 1, further comprising:
randomly selecting a predetermined number of
30 documents from a remaining subset of the collection of documents not in the responsive documents subset; and

determining whether the randomly selected documents include additional relevant documents.

6. The method of claim 5, further comprising:
5 identifying one or more specific terms in the additional relevant documents that render the documents relevant;
expanding the query terms with the specific terms; and
10 re-running at least the search with the expanded query terms.

7. The method of claim 5, wherein if the randomly selected documents includes one or more
15 additional relevant documents, the query terms are expanded and the search is re-run with the expanded query terms.

8. The method of claim 7, further comprising
20 comparing a ratio of the additional relevant documents and the randomly selected documents to a predetermined acceptance level, to determine whether to apply a refined set of query terms.

25 9. The method of claim 1, further comprising:
selecting two or more search terms;
identifying synonyms of the search terms; and
forming the query terms based on the search terms and synonyms.

30

10. The method of claim 1, further comprising:

identifying a correspondence between a sender and a recipient, in the responsive documents subset;

automatically determining one or more additional documents which are in a thread of the correspondence, the additional documents not being in the responsive documents subset; and

adding the additional documents to the responsive documents subset.

10 11. The method of claim 1, further comprising:
determining whether any of the documents in the responsive documents subset includes an attachment that is not in the responsive documents subset; and
adding the attachment to the responsive documents
15 subset.

12. The method of claim 1, further comprising applying a statistical technique to determine whether remaining documents not in the responsive documents set meets a predetermined acceptance level.

13. The method of claim 1, wherein the search includes (a) a Boolean search of the collection of documents based on the plurality of query terms, the Boolean search returning a first subset of responsive documents from the collection, and (b) a second search by applying a recall query based on the plurality of query terms to remaining ones of the collection of documents which were not returned by the Boolean search, the second search returning a second subset of responsive documents in the collection, and wherein

the responsive documents subset is constituted by the first and second subsets.

14. The method of claim 13, wherein the first
5 Boolean search applies a measurable precision query based on the plurality of query terms.

15. The method of claim 14, further comprising:
automatically tagging each document in the first
10 subset with a precision tag;
reviewing the document bearing the precision tag to determine whether the document is properly tagged with the precision tag; and
determining whether to narrow the precision query
15 and rerun the Boolean search with the narrowed query terms.

16. The method of claim 13, further comprising:
automatically tagging each document in the second
20 subset with a recall tag;
reviewing the document bearing the recall tag to determine whether the document is properly tagged with the recall tag; and
determining whether to narrow the recall query
25 and rerun the second search with the narrowed query terms.

17. The method of claim 13, further comprising
reviewing the first and second subsets to determine
30 whether to modify the query terms and rerun the Boolean search and second search with modified query terms.

18. A method for reviewing a collection of documents to identify relevant documents from the collection, the method comprising:

5 running a search of the collection of documents, based on a plurality of query terms, the search returning a subset of responsive documents in the collection;

10 automatically identifying a correspondence between a sender and a recipient, in the responsive documents subset;

15 automatically determining one or more additional documents which are in a thread of the correspondence, the additional documents not being in the responsive documents subset; and

adding the additional documents to the responsive documents subset.

19. The method of claim 18, wherein the search is applied through a search engine.

20. The method of claim 18, wherein the search includes a concept search, and the concept search is applied through a concept search engine.

25

21. The method of claim 18, further comprising: determining for each document in the responsive documents subset, a corresponding probability of relevancy; and

30 removing from the responsive documents subset documents that do not reach a threshold probability of relevancy.

22. The method of claim 21, wherein the probability of relevancy of a document is scaled according to a measure of obscurity of the search
5 terms found in the document.

23. The method of claim 18, further comprising applying a statistical technique to determine whether a remaining subset of the collection of documents not
10 in the responsive documents subset meets a predetermined acceptance level.

24. The method of claim 18, further comprising:
randomly selecting a predetermined number of
15 documents from a remainder of the collection of documents not in the responsive documents subset;
determining whether the randomly selected documents include additional relevant documents;
identifying one or more specific terms in the
20 additional relevant documents that render the documents relevant;
expanding the query terms with the specific terms; and
re-running the search with the expanded query
25 terms.

25. The method of claim 18, further comprising:
randomly selecting a predetermined number of documents from a remainder of the collection of
30 documents not in the responsive documents subset;
determining whether the randomly selected documents include additional relevant documents;

comparing a ratio of the additional relevant documents and the randomly selected documents to a predetermined acceptance level; and

expanding the query terms and rerunning the
5 search with the expanded query terms, if the ratio does not meet the predetermined acceptance level.

26. The method of claim 18, further comprising:
selecting two or more search terms;
10 identifying synonyms of the search terms; and
forming the query terms based on the search terms and synonyms.

27. The method of claim 18, further comprising:
15 determining whether any of the responsive documents in the responsive documents subset includes an attachment that is not in the subset; and
adding the attachment to the subset.

20 28. A method for reviewing a collection of documents to identify relevant documents from the collection, the method comprising:

running a search of the collection of documents, based on a plurality of query terms, the search
25 returning a subset of responsive documents in the collection;

automatically determining whether any of the responsive documents in the responsive documents subset includes an attachment that is not in the
30 subset; and

adding the attachment to the responsive documents subset.

29. The method of claim 28, wherein the search is applied through a search engine.

5 30. The method of claim 28, wherein the search includes a concept search, and the concept search is applied through a concept search engine.

31. The method of claim 28, further comprising:
10 determining for each document in the responsive documents subset, a corresponding probability of relevancy; and
removing from the responsive documents subset documents that do not reach a threshold probability of
15 relevancy.

32. The method of claim 31, wherein the probability of relevancy of a document is scaled according to a measure of obscurity of the search
20 terms found in the document.

33. The method of claim 28, further comprising applying a statistical technique to determine whether a remaining subset of the collection of documents not
25 in the responsive documents subset meets a predetermined acceptance level.

34. The method of claim 28, further comprising:
randomly selecting a predetermined number of
30 documents from a remainder of the collection of documents not in the responsive documents subset;

determining whether the randomly selected documents include additional relevant documents;

identifying one or more specific terms in the additional responsive documents that render the documents relevant;

expanding the query terms with the specific terms; and

re-running the search with the expanded query terms.

10

35. The method of claim 28, further comprising:

selecting two or more search terms;

identifying synonyms of the search terms; and

forming the query terms based on the search terms

15 and synonyms.

36. The method of claim 28, further comprising:

identifying a correspondence between a sender and a recipient, in the responsive documents subset;

20 automatically determining one or more additional documents which are in a thread of the correspondence, the additional documents not being in the responsive documents subset; and

25 adding the additional documents to the responsive documents subset.

37. A method for reviewing a collection of documents to identify relevant documents from the collection, the method comprising:

30 running a search of the collection of documents, based on a plurality of query terms, the search

returning a subset of responsive documents from the collection;

randomly selecting a predetermined number of documents from a remainder of the collection of documents not in the responsive documents subset;

determining whether the randomly selected documents include additional relevant documents;

identifying one or more specific terms in the additional responsive documents that render the documents relevant;

expanding the query terms with the specific terms; and

re-running the search with the expanded query terms.

15

1/6

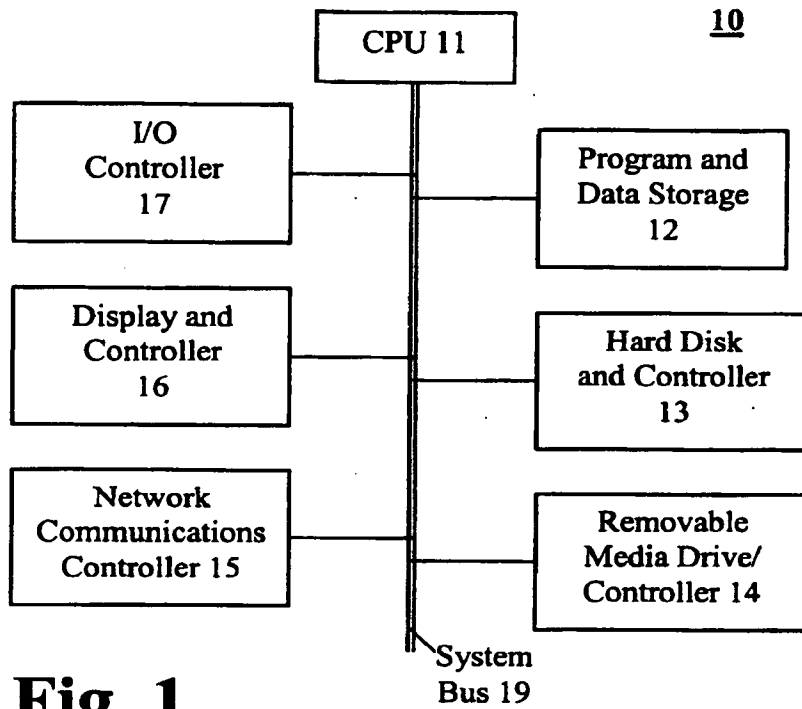


Fig. 1

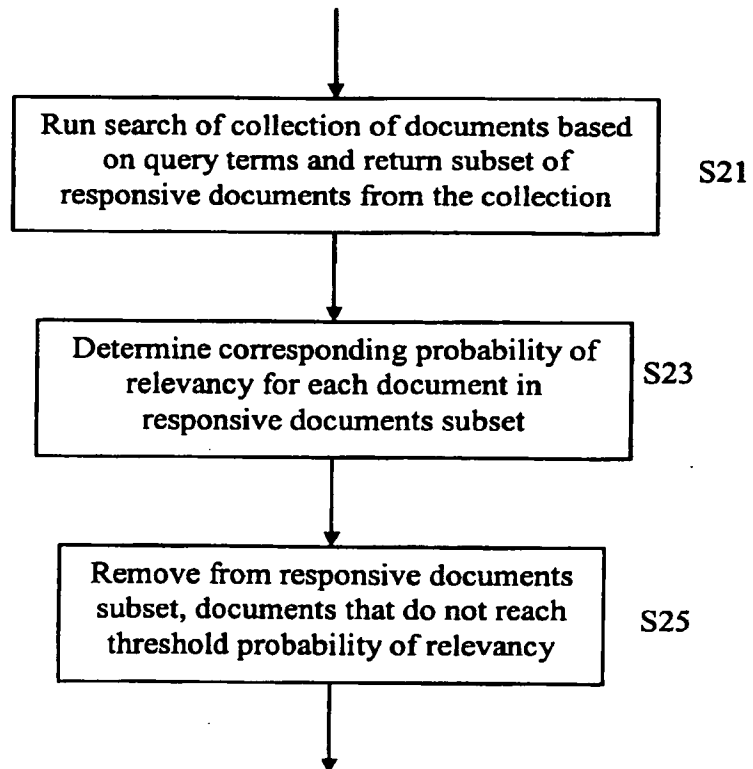


Fig. 2

2/6

Fig. 3

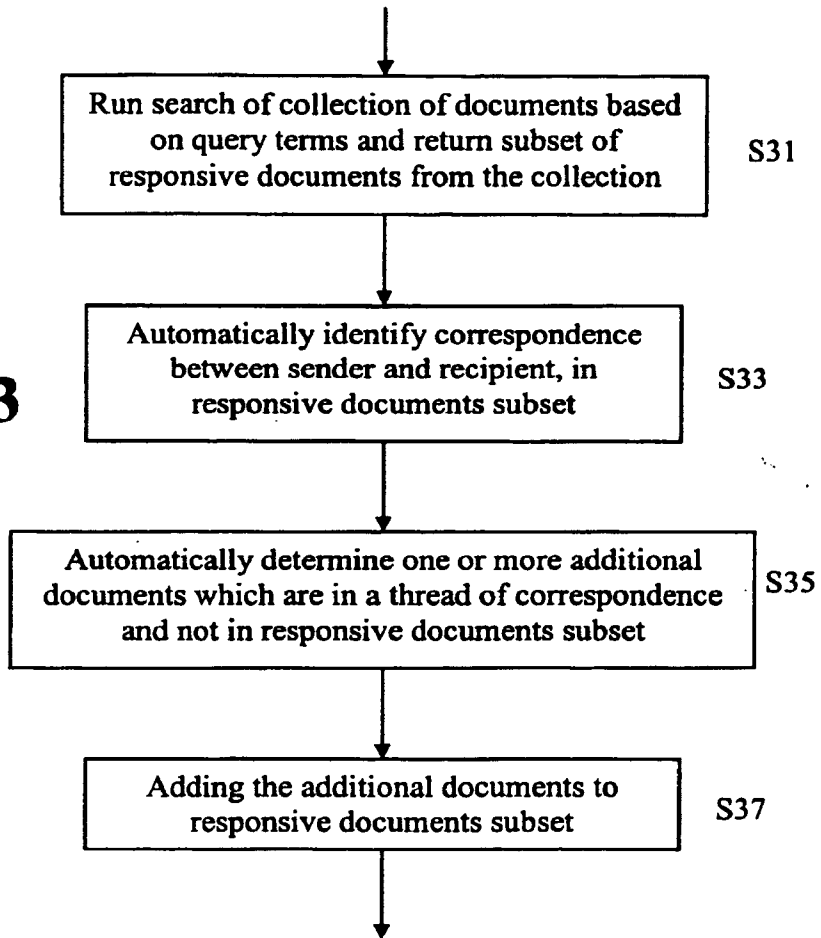
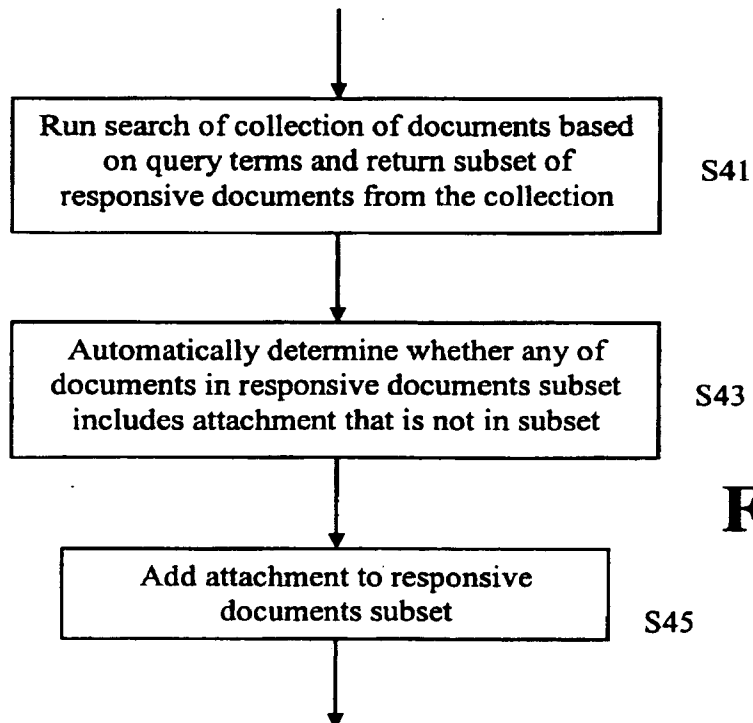
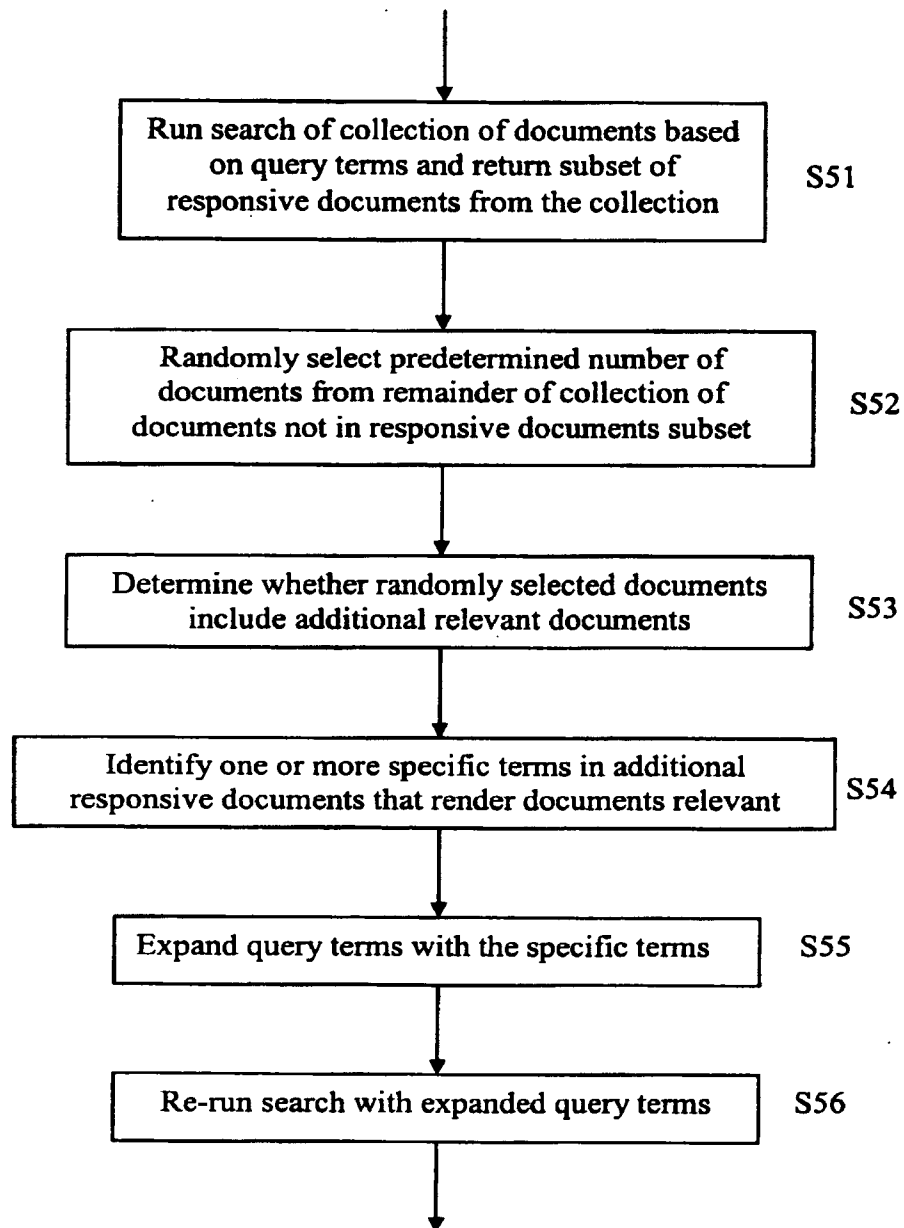


Fig. 4



3/6

**Fig. 5**

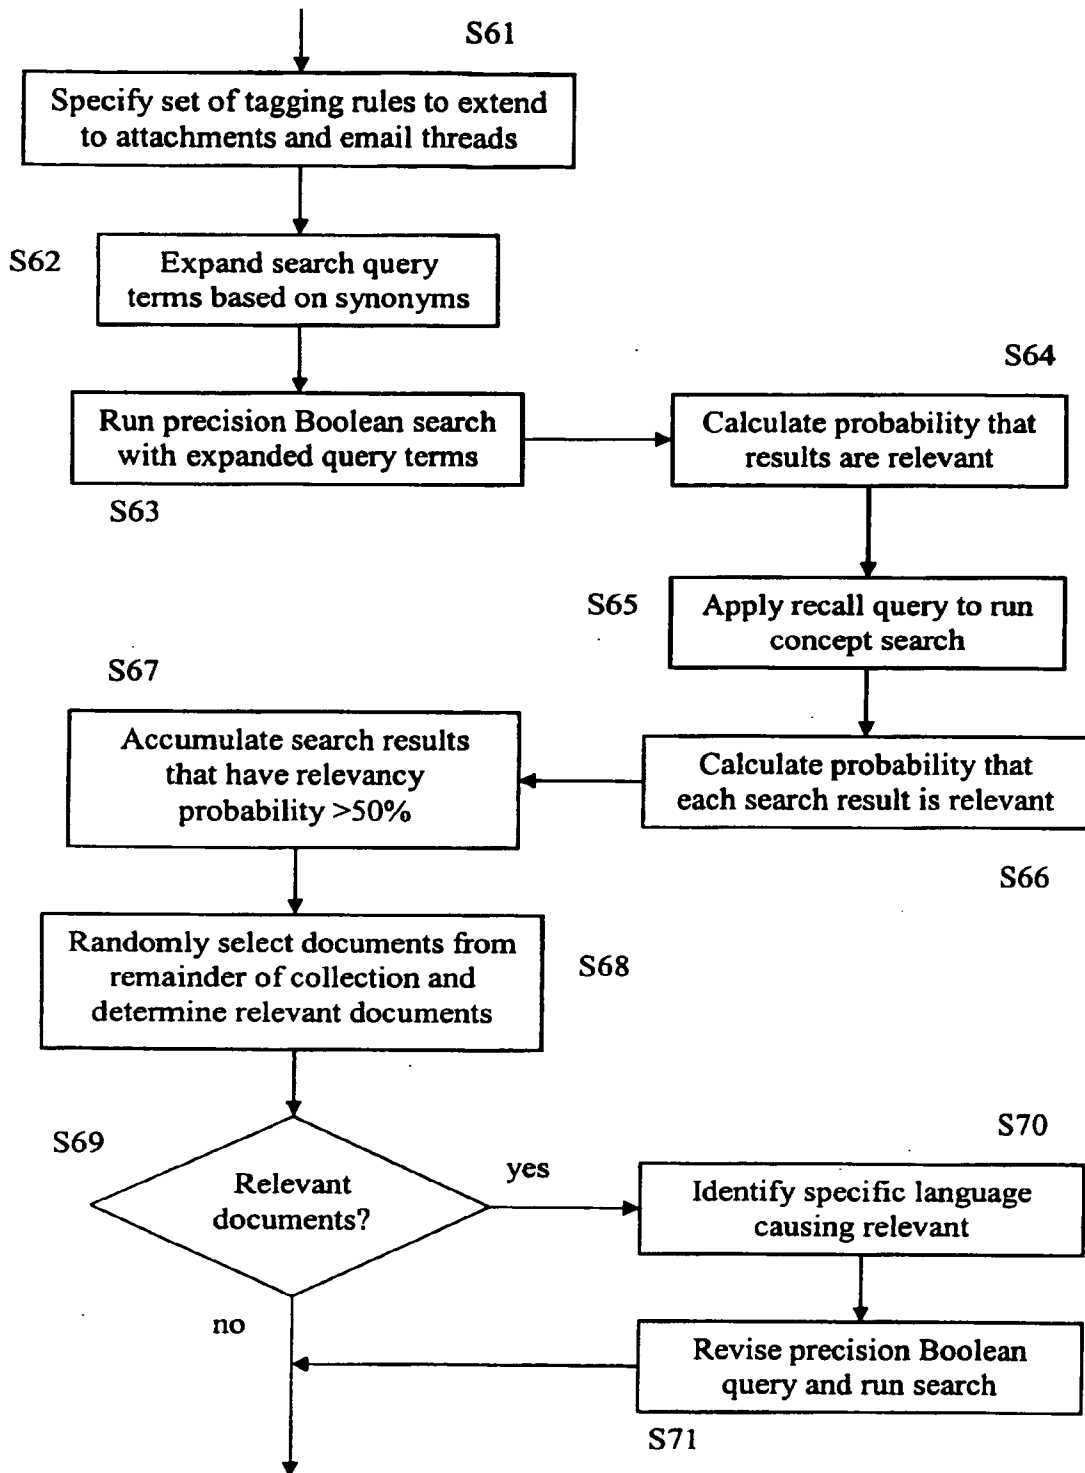


Fig. 6

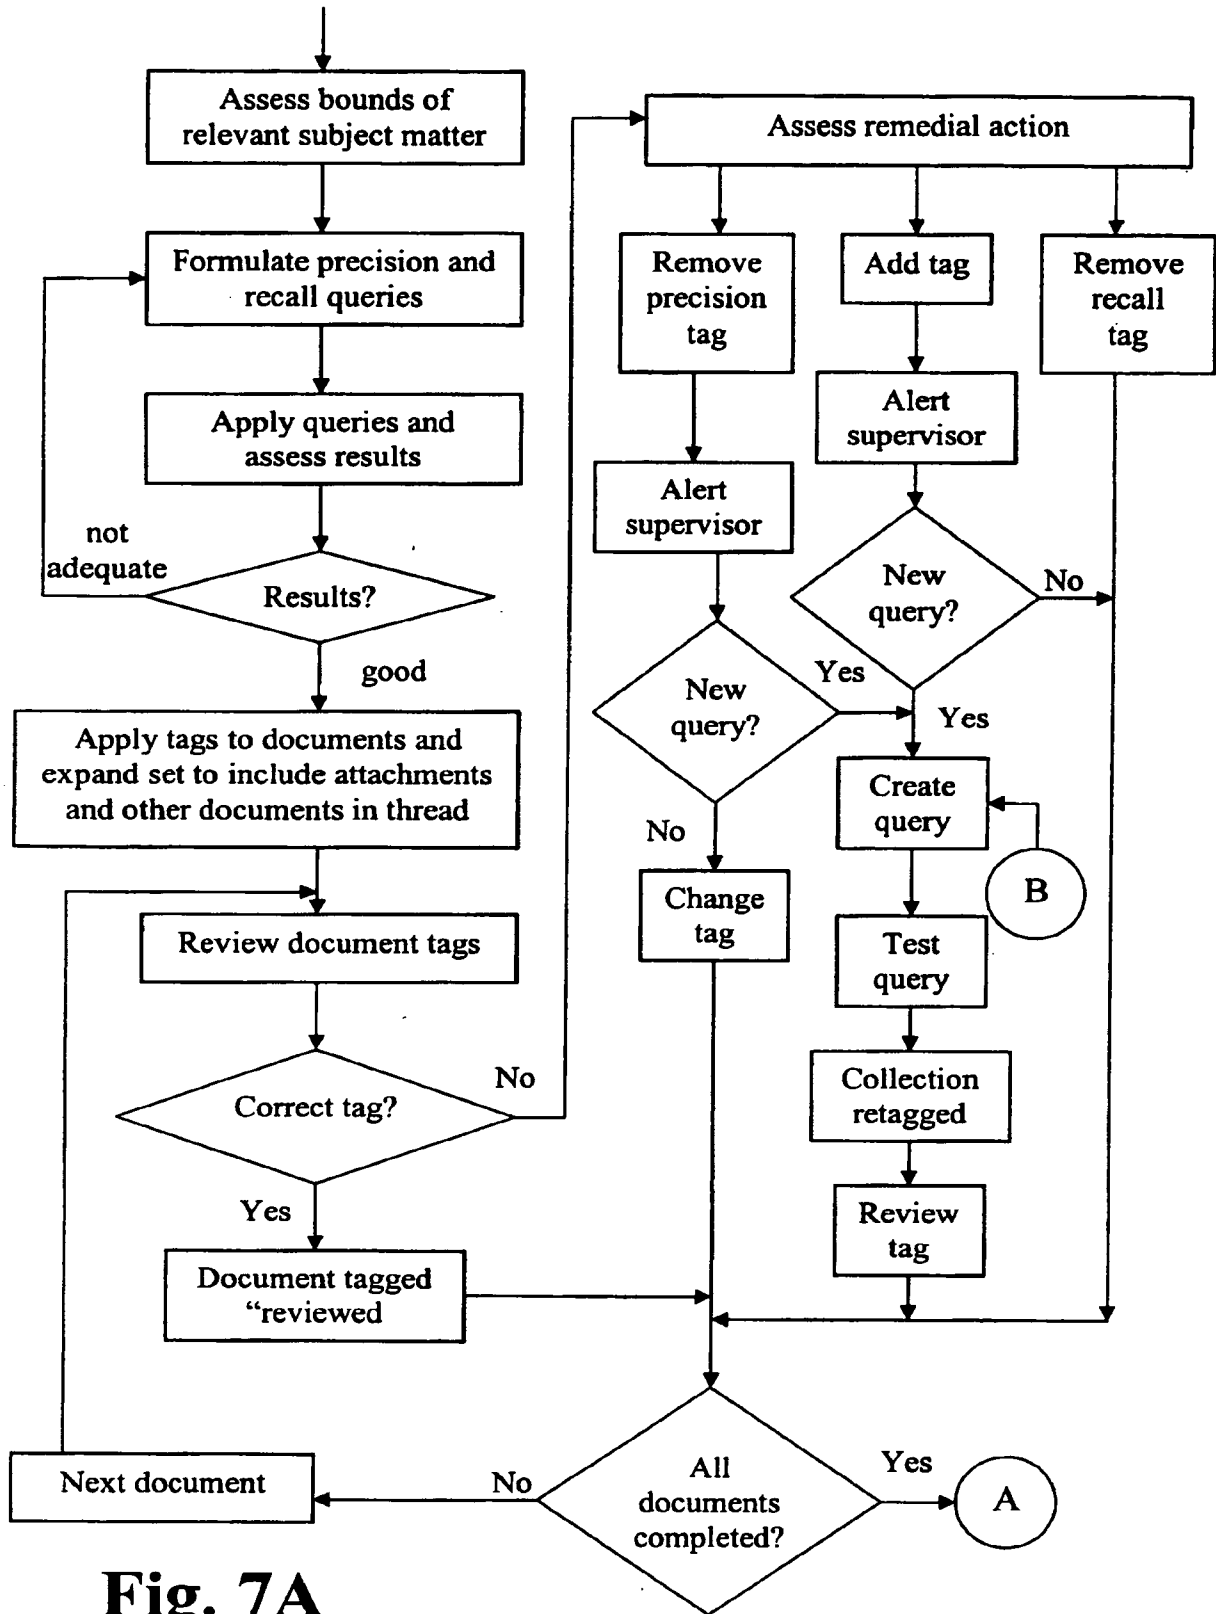


Fig. 7A

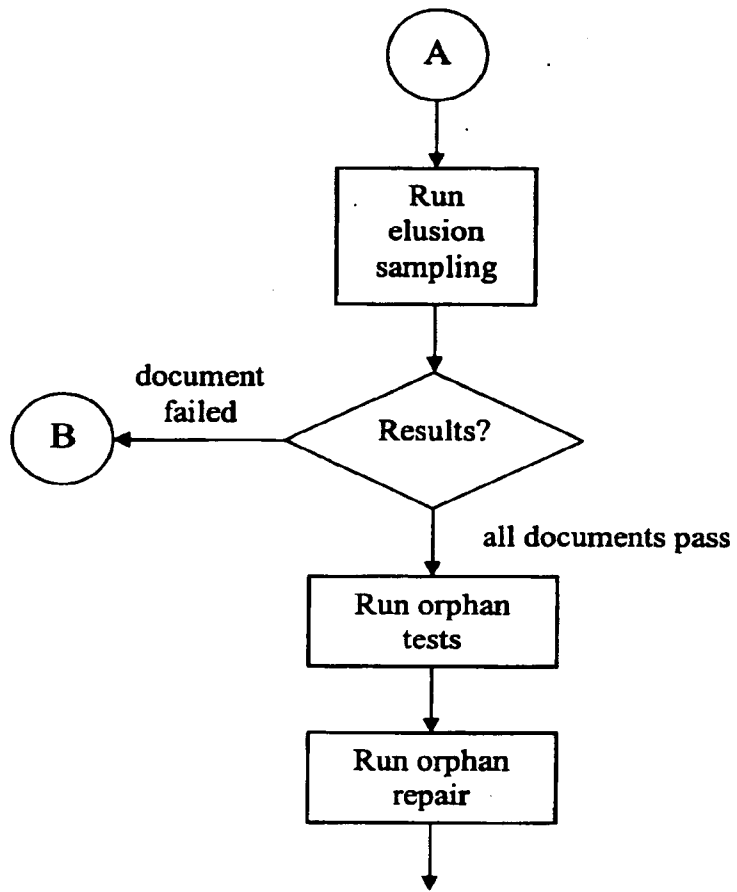


Fig. 7B