

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
21 February 2008 (21.02.2008)

PCT

(10) International Publication Number
WO 2008/022150 A2

(51) International Patent Classification:
G06F 19/00 (2006.01)

(74) Agents: **FORD, Stephen, S.** et al.; Marger Johnson & McCollom PC, 210 SW Morrison St. Suite 400, Portland, OR 97204 (US).

(21) International Application Number:
PCT/US2007/075929

(22) International Filing Date: 14 August 2007 (14.08.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/464,443 14 August 2006 (14.08.2006) US
11/464,446 14 August 2006 (14.08.2006) US

(71) Applicant (for all designated States except US): **IN-QUIRA, INC.** [US/US]; 851 Traeger Avenue, Suite 125, San Bruno, CA 94066 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **COOPER, Edwin, Riley** [US/US]; 6521 Dover Street, Oakland, CA 94609 (US). **DUKES, Michael, Peter** [NZ/US]; 1041 Iris Avenue, Sunnyvale, CA 94086 (US). **BIERNER, Gann, Alexander** [US/US]; 851 Traeger Avenue, Suite 125, San Bruno, CA 94066 (US). **BEGHELLI, Filippo, Ferdinando, Paulo** [US/GB]; 40 Swallow Court, Admiral Way, London W9 3TX (GB).

(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).

Published:

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

(54) Title: METHOD AND APPARATUS FOR IDENTIFYING AND CLASSIFYING QUERY INTENT

(57) Abstract: Linguistic analysis is used to identify queries that use different natural language formations to request similar information. Common intent categories are identified for the queries requesting similar information. Intent responses can then be provided that are associated with the identified intent categories. An intent management tool can be used for identifying new intent categories, identifying obsolete intent categories, or refining existing intent categories.



WO 2008/022150 A2

METHOD AND APPARATUS FOR IDENTIFYING AND CLASSIFYING QUERY INTENT

Cross-Reference to Related Applications

The present application claims priority from U.S. Patent Application No. 11/464,443 filed on August 14, 2006 and U.S. Patent Application No. 11/464,446 filed on August 14, 2006, which are incorporated herein by reference in their entirety for all purposes.

Background

Search engines try to provide the most relevant responses to user questions. Unfortunately, many search engines return information that may be unrelated, or not directly related, to the question. For example, search engines may return any document containing words matching keywords in the question. The user has to then manually sort through each returned document in an attempt to identify information that may be relevant or answer the question. This “brute force” method is time consuming and often fails to locate the precise information sought in the question.

Current search engines try to help the user in their manual document search by ranking the returned documents. This ranking method may rank documents simply according to the number of words in the documents that match keywords in the query. At least one critical limitation with this keyword search technique is that the user may not necessarily input search terms needed by the search engine to locate the correct information. In addition, even appropriate keywords may also be associated with other documents unrelated to the information sought by the user.

Search engines have been developed that attempt to classify queries. For example, the search engine may try to associate different words in the search query with different

information categories. The search engine then attempts to provide the user with responses associated with the identified information category.

A critical problem with these information retrieval schemes is that there are seemingly limitless ways in a natural language for a user to request for the same information. And as
5 also mentioned above, the user may not necessarily enter, or even know, the best words or phrases for identifying the appropriate information. Accordingly, the search engine can only classify a very limited number of subject matters. Further, a large amount of human resources are required to keep these types of search engines up to date with new information categories that may develop over time. Thus, these “higher level” search engines have had
10 only limited success providing responses to user questions.

The present invention addresses this and other problems associated with the prior art.

Summary of the Invention

Linguistic analysis is used to identify queries that use different natural language
15 formations to request similar information. Common intent categories are identified for the queries requesting similar information. Intent responses can then be provided that are associated with the identified intent categories. An intent management tool can be used for identifying new intent categories, identifying obsolete intent categories, or refining existing intent categories.

20 The foregoing and other objects, features and advantages of the invention will become more readily apparent from the following detailed description of a preferred embodiment of the invention which proceeds with reference to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a graph that shows how the number of unique queries received by an enterprise can be reduced by classifying the queries into intent categories.

FIG. 2 is a block diagram showing an intent based search engine.

FIG. 3A is a graph showing how the intent based search engine can provide different
5 types of responses according to query frequency.

FIG. 3B is a block diagram showing the intent based query engine in more detail.

FIG. 4 is a block diagram showing how an intent management tool is used for managing intent categories.

FIG. 5A is a block diagram showing how intent categories are automatically
10 identified and derived from queries.

FIG. 5B shows in more detail how intent categories are automatically identified and derived from queries.

FIG. 6 is a flow diagram showing in more detail how the intent management tool in FIG. 4 manages intent categories.

15 FIG. 7 is a flow diagram showing how new intent categories can be generated and refined using the intent management tool.

FIG. 8 shows how intent categories can be associated with an intent hierarchy.

FIG. 9 is a block diagram showing how different intent responses are displayed according to an associated intent hierarchy.

20 FIG. 10 is a block diagram showing how a user or administrator can associate different parameters with an intent category.

FIG. 11 is a flow diagram showing how user parameters can be associated with an intent category.

FIG. 12 is a block diagram showing how features from an ontology are used to
25 identify query clusters.

FIG. 13 is a flow diagram showing how clustering is used to generate new intent categories.

FIG. 14 is a block diagram showing how new intent categories identified in FIG. 13 are associated with different locations in an intent hierarchy.

5 FIG. 15 is a block diagram showing how parameters can be assigned to intent responses.

DETAILED DESCRIPTION

FIG. 1 is a graph showing the relationship between unique search queries and the
10 frequency they are received by a particular enterprise. The enterprise could be any type of business or institution, such as a car dealer, financial service, telecommunications company, etc. that has a search engine that attempts to provide responses to questions submitted by users via computer terminals. The horizontal axis 12 refers to the different unique queries that may be received by one or more search engines operated by the enterprise. The vertical
15 axis 14 refers to the frequency of the different unique queries.

A challenge exists trying to algorithmically decipher the meaning of received queries and then provide responses associated with the identified query meaning. For example, the shape of curve 16 indicates that a first portion 20 of the unique queries occur with the most frequency and a second portion 18 of the queries occur with slightly less frequency. As can
20 be seen, a large percentage of the total number of queries occur in this second lower frequency portion 18.

Due to maintenance and resource issues, it may only be possible for search engines to try and determine the meaning and then provide associated responses for a subset of the most frequency received queries. For example, a search engine for an online book retailer may be
25 designed to look for and identify queries related to book prices. However, it may not be cost

and time effective to design the search engine to try and determine the meaning and provide associated responses for every possibly book question. For example, thousands of possible responses would have to be configured just to capture a relatively small percentage of possible book questions. This large number of preconfigured responses are difficult to maintain and would have to be constantly updated to respond to the never ending number of new questions related to new books.

Unfortunately and according to curve 16, developing a search engine that is only capable of responding to the most frequency asked questions 20, ignores a large percentage of queries 18 that may be received by the online book retailer. This substantial portion of “query outliers” 18 would then have to be processed using conventional keyword searches. The limitations of key word searching was previously explained above.

As a result, the search engine for the online book retailer may not provide the most relevant responses for a large percentage of the received queries. This could negatively effect business. In the online book seller example, some of the less frequently received queries 18 may relate to rare books that may have larger mark ups than the more recent/popular books associated with the most frequently received queries 20. Accordingly, these search engine limitations may cause the online book retailer to lose some high profit rare books sales.

An intent based search engine is used to determine the intent categories of queries and then provide corresponding responses for a larger proportion of unique queries that may be received by an enterprise. “Intent” refers to the meaning associated with a query. An intent based search engine classifies multiple different unique queries into common “useful” intent categories. The term “useful” refers to intent categories that are associated with relevant responses to information requests. For example, identifying an intent category for a group of queries associated with “the Internet”, may be too broad to be useful to an enterprise that is attempting to respond to queries related to vehicle sales. However, identifying an intent

category associated to “purchasing a vehicle over the Internet”, may be very useful when responding to a user query.

Classifying queries according to their intent category changes the relationship between unique queries 12 and their frequency 14. This is represented by curve 22 where a significantly larger portion 20 of all received queries can be classified by a relatively small number of intent categories. Many of the outlier queries 18 previously located underneath curve 16 can be identified as having the same meaning or “intent” as some of the more frequently asked queries 20. Identifying the intent category for queries allow the search engine to provide more relevant responses to a larger percentage of queries while at the same time requiring substantially fewer resources to maintain the search engine. In other words, fewer responses can be used to adequately respond to a larger number of queries.

For example, a large number of queries received by a financial services enterprise may be related to 401K retirement plans. The queries related to 401Ks may be expressed by users in many different ways. For instance, “what is the current value of my 401K”, “how much money is in my company retirement account”, “show me the status of my 401K investments”, etc. The information sought for each of these queries can be classified by the same intent category, namely: Intent Category = value of 401K. By classifying queries into intent categories, fewer associated responses have to be maintained.

FIG. 2 shows a computer network system 30 that includes an enterprise 32 that has one or more enterprise servers 34 and one or more enterprise databases 36. As described above, the enterprise 32 may be an online retailer that sells books and other retail items. In this example, the enterprise database 36 may contain price lists and other information for all of the books and other merchandise available for purchase. In another example, the enterprise 32 may be associated with a car dealership or financial institution and the enterprise database 36 could include vehicle or financial information, respectively. These

are, of course, just examples, and any type of business or entity can be represented as enterprise 32.

Other web servers 26 may operate outside of the enterprise 32 and may include associated files or other web content 28. Examples of content stored in enterprise database 36 and in web server 26 may include HyperText Markup Language (HTML) web pages, Portable Document Format (PDF) files, Word® documents, structured database information or any other type of electronic content that can contain essentially any type of information.

Information in database 36 may be stored in a structured preconfigured format specified for the enterprise 32. For example, a book or vehicle price list may be considered structured content. The enterprise 32 may also generate and store specific intent responses either in enterprise database 36 or on enterprise server 34 that are associated with specific intent categories 50. Other information that is contained in enterprise database 36, or contained on other web servers 26, may be considered non-structured content. This may include HTML web pages, text documents, or any other type of free flowing text or data that is not organized in a preconfigured data format.

A query 46 (e.g., electronic text question) may be initiated by a user from a terminal 25 through a User Interface (UI) 40. The terminal 25 in one example may be a Personal Computer (PC), laptop computer, wireless Personal Digital Assistant (PDA), cellular telephone, or any other wired or wireless device that can access and display content over a packet switched network. In this example, the query 46 is initiated from the UI 40 and transported over the Internet 48 to the enterprise server 34. For example, query 46 may be a question sent to a bank asking: Query="what is the current interest rates for CDs".

The enterprise server 34 operates a novel intent based search engine 35 alternatively referred to as an automated response system. The search engine 35 provides electronic responses, answers, and/or content pursuant to electronically submitted queries 46. The

intent based search engine 35 uses a set of predetermined intent categories 50, one or more ontologies 52, and an Intelligent Matching Language (IML) engine 53 to identify the intent category 51 for query 46 and then provide an associated intent response 49.

The intent analysis is described in more detail below and converts the relatively flat query vs. frequency relationship curve 16 previously shown in FIG. 1 into the steeper query intent vs. frequency relationship curve 22. This results in the intent based search engine 35 presenting a more relevant intent based response 44 for electronically submitted question 42 while at the same time requiring a relatively low number of intent responses 49 for responding to a large number of unique queries 46. Accordingly, fewer resources have to be maintained by the intent based search engine 35.

The search engine 35 receives queries 46 from the UI 40 resulting from a question 42 entered by a user. The search engine 35 attempts to match the meaning or "intent" of the query 46 with preconfigured intent categories 50 using an intelligent matching language engine 53 and ontologies 52. The intent based search engine 35 then identifies one of the intent based responses 49 associated with the identified query intent category 51. The intent responses 49 may be preconfigured content or network links to information responsive to associated intent categories 50. The intent responses 49 can also include any structured and/or non-structured content in the enterprise database 36 or on different web servers 26 that the intent based search engine 35 associates with the identified intent category 51. The identified information is then sent back to the UI 40 as intent based response 44.

The enterprise server 34 can include one or more processors that are configured to operate the intent based search engine 35. The operations performed by the intent based search engine 35 could be provided by software computer instructions that are stored in a computer readable medium, such as memory on server 34. The instructions are then executed by one or more of the processors in enterprise server 34. It should also be understood that the

examples presented below are used for illustrative purposes only and the scope of the invention is not limited to any of the specific examples described below.

In one embodiment, the intent categories 50 are represented using a natural language, such as used in the IML engine 53. Using a natural language allows a system administrator to more easily create, delete, and modify intent categories 50. For example, the administrator can more easily identify which characteristics in an intent category need to be changed to more effectively classify a group of queries with a particular intent category. This is more intuitive than present information retrieval systems that use statistical analysis to classify queries into different categories.

Referring to FIGS. 2 and 3A, some of the different operations are described that may be performed by the intent based search engine 35. The search engine 35 may identify “a priori”, the most frequently queried intent categories 50 (FIG. 2), and automatically display associated intent responses on the enterprise website. For example, 5-10% of the queries received by a financial service enterprise may contain questions related to retirement accounts. Accordingly, static links to web pages containing retirement account information may be presented on the home webpage for the financial institution prior to the user ever entering a question. This is referred to as pre-query based responses 60.

The intent based search engine 35 provides other types of responses according to the type of information that can be derived from the received queries. For example, there may be a set of around 100 intent categories 50 and corresponding intent responses 49 that address 60-70% of all unique queries received by a particular enterprise 32 (FIG. 2). A set of intent categories 50 are either manually or automatically derived based on previously received query information that cover this large percentage of unique queries. A set of intent responses 60 or 62 are then created that respond to the most frequently queried intent categories 50. The search engine 35 then attempts to match received queries with one of

these frequent intent categories 50 and, if successful, sends back the corresponding intent based responses 60 or 62 (FIG. 3A).

Any identified intent categories 50 can also be used to improve the relevance of any other information provided in response to the query. For example, the identified intent category 50 in FIG. 2 may be used to identify both preconfigured intent based responses 49 and/or used for conducting an additional document search for other content in enterprise database 36 (FIG. 2) or other web content 28 in other web servers 26. The identified intent category 50 can also be used to extend or limit the scope of a document search or used to change the rankings for documents received back from the search.

The search engine 35 may use the IML engine 53 and ontologies 52 to discover concepts and other enterprise specific information contained in the queries 64. Any concepts discovered during query analysis can be used to discover the intent categories and associated intent based responses 62. However, the search engine 35 may not be able to identify an intent category 50 for some percentage of less frequently received queries. If no intent categories can be identified, the search engine 35 can still use any identified concepts to provide ontology based responses 64.

To explain further, the ontologies 52 shown in FIG. 2 may associate different words such as IRA, 401K, Roth, retirement, etc., with the broader concept of retirement accounts. The Intelligent Matching Language (IML) engine 53 is used in combination with the ontologies 52 to identify and associate these different phrases, words, and word forms, such as nouns, adjectives, verbs, singular, plural, etc., in the queries with different concepts.

For example, the IML engine 53 may receive a question asking about the price of a book but that does not necessarily include the symbol "\$", or use the word "dollar". The IML engine 53 may use ontologies 52 to associate the symbol "\$" and the word "dollar" with the words "Euro", "bucks", "cost", "price", "Yen", etc. The IML engine 53 then applies

concepts such as <dollar> or <price> to the query 46 to then identify any words in the query associated with the <dollar> or <price> concepts. The identified concepts, words, etc., identified using IML engine 53 and ontology 52 are then used by the intent based search engine 35 to search for a relevant response.

5 One example operation of an IML engine 53 is described in co-pending patent application Serial No. 10/820,341, filed April 7, 2004, entitled: AN IMPROVED ONTOLOGY FOR USE WITH A SYSTEM, METHOD, AND COMPUTER READABLE MEDIUM FOR RETRIEVING INFORMATION AND RESPONSE TO A QUERY, which is herein incorporated by reference.

10 The intent based search engine 35 may also combine conventional keyword analysis with other intent and ontology based query analysis. If the search engine 35 does not identify any intent categories or ontology based concepts with the query, keyword based responses 66 may be provided based solely on keyword matching. For example, any content containing the same keywords used in the query 46 can be provided to the UI 40.

15 Thus, when no intent category can be determined, the search engine 35 may still use the domain based knowledge from the ontologies 52 to discover the most relevant responses 64. Alternatively, when the domain based knowledge does not provide any further understanding as to the meaning of the query, keyword analysis is used to provide keyword based responses 66. This is, of course, just one example of different combinations of intent,
20 ontology concepts, and keyword analysis that can be performed on a query to provide different intent based responses 62, ontology based responses 64, and keyword based responses 66.

The intent based search engine 35 may conduct all of this intent, ontology and keyword analysis at the same time and then provide responses based on the highest level of
25 query understanding. The intent based search engine 35 can use any other type of word,

phrase, sentence, or other linguistic analysis, to determine the intent, concepts and words in the query. Similarly, any number of intent categories 50 and intent responses 49 may be used by the intent based search engine 35 and may cover any percentage of the unique queries 60, 62, 64, and 66 received by the enterprise server 34.

5 As described in more detail below, the intent based search engine 35 allows more efficient administration of an enterprise information system. For example, the most frequently referred to intent categories can be identified and associated intent responses derived. This provides a substantial advantage over existing search engine administration where little or no ability exists for classifying multiple different queries with the same
10 associated response. Similarly, the administrator can more efficiently add, delete, update, and/or modify the most relevant intent categories. In other words, the administrator is less likely to waste time generating or maintaining responses for infrequently received or irrelevant queries. This again is described in more detail below.

FIG. 3B shows in more detail how the search engine 35 identifies an intent category
15 51 for a query 46 and then provides an associated intent response 44. The query 46 received by the intent based search engine 35 is first analyzed by IML engine 53. The IML engine 53 uses natural language linguistic analysis to match the query 46 with one of the intent categories 50. One or more ontologies 52 are used that associated different words, phrases, etc., with different concepts that may be industry specific for the enterprise. For example,
20 login, permission, password, passcode, etc., may all be associated with an <account information> concept that is regularly referred to by users accessing the enterprise network.

The IML 53 uses the ontologies 52, as well as other natural language associations, when trying to match the query 46 with one of the intent categories 50. When the query 46 is matched with one of the intent categories 50, the search engine 35 identifies the intent

response 44 associated with the identified intent category 51. The identified intent response 44 is then displayed to the user.

For example, the following queries may either be received by the intent based search engine 35 in FIG. 2.

5	<u>Queries</u>	<u>Intent Category</u>
	How do I change my password?	
	I want to change my password	
	How do I update my pass word?	
	Is there a form for modifying passcodes?	Change Password
10	Change my password	
	Need to change my secret code	
	Is there a password change form?	

As seen above, each of these queries is associated with the same “change password” intent category 51. The search engine 35 may match each of these questions with the same intent category 51 and provide a corresponding intent response 44.

Intent Management Tool

FIG. 4 shows an intent management tool 67 that can be used to identify the most frequently queried intent categories 69A, identify the least frequently queried intent categories 69B, generate new intent categories 69C, identify queries 69D that do not match any existing intent categories, generate intent category hierarchies 69E, and/or assign and identify parameters to intent categories 69F.

The intent management tool 67 receives queries 68 that have been logged for some period of time by the enterprise server 34 (FIG. 2). The intent management tool 67 then uses existing intent categories 50, the intelligent matching language engine 53, and ontologies 52 to identify different information related to the logged queries 68 and the intent categories 50. For example, it may be beneficial to a website administrator to know which intent categories 69A match the most logged queries 68 or which intent categories 69B match the fewest

logged queries 68. The intent management tool 67 can also automatically generate a new intent category 69C or identify queries 69D that do not match any existing intent categories. An intent category hierarchy 69E can be created that is used for providing alternative responses to received queries. Multiple different parameters 69F can also be identified and
5 assigned to different intent categories and then used for associating the intent categories with different intent responses.

FIG. 5A shows one example of how the intent management tool 67 generates a new intent category 79, or how the intent based search engine 35 matches a query with an existing intent category 79. Multiple different queries 70 may be received or logged by the enterprise
10 server 34 (FIG. 2). Some of these queries 70A-70E may be associated with a same existing or non-existing intent category: CONTRACT EXPIRATION SUPPORT. A first one of the queries 70A however contains the question: “when idoes, my service expire”. In a first spelling analysis stage 72, a natural language engine 71 used either by the intent based search engine 35 (FIG. 2) or intent management tool 67 (FIG. 4) checks the spelling in query 70A.
15 The term “idoes” is not found in any of the ontologies 52 (FIG. 2). Accordingly, “idoes” is replaced with the closest match “does”.

In a next punctuation and capitalization stage 73, punctuation is analyzed and a comma is removed that does not make sense. A speech analysis stage 74 analyzes nouns, verbs, etc., in the corrected question to generate an annotated question. For example, the
20 word “does” is identified as a verb and the word “service” is identified as a noun phrase. Other words in an identified noun phrase may be removed. In a stem analysis stage 75, the search engine 35 or management tool 67 may add other forms of the same identified words to the annotated question. For example, the identified noun “service” could also be associated with “servicing”, “serviced”, etc.

In a concept analysis stage 76, the management tool uses the ontologies 52 to identify any concepts associated with the annotated question. For example, the word “service” may be associated with the concept <phone service contract> and the word “expire” may be associated with the concept <end>. A linguistic analysis stage 77 then adds linguistic analysis results. For example, the annotated question may be determined to be related to an account, and accordingly restricted to related account and marketing information. In an intent analysis stage 78, an existing intent category is identified that matches the annotated question or a new intent category 79 is created for the previously annotated question. Similar linguistic analysis of questions 70B-70E may result in identifying the same existing intent category 79 or may be used along with query 70A to create a new intent category 79.

FIG. 5B describes in more detail how the intent categories are identified and created. A query 70F may ask the question: “I’m having trouble with my cell phone”. The natural language engine 71 in combination one or more ontologies 52 are then used to conduct the concept analysis 76 and linguistic analysis 77 previously described in FIG. 5A. Different ontologies 52A-52C can be associated with different concepts. For example, ontology 52A is associated with the concept <trouble>, the ontology 52B is associated with pronouns and adjectives related to the concept <my>, and ontology 52C is associated with nouns and noun phrases related to the concept <cell phone>.

The natural language engine 71 uses ontologies 52 to identify different concepts 81A, 81B and 81C associated with the query 70F. The natural language engine 71 may identify a <my> concept 81A and a <trouble> concept 81C in query 70F. The natural language engine 71 may also identify the first <my> concept 81A as preceding a noun phrase 81B and also being located in the same sentence as the <trouble> concept 81C. This combination of concepts and associated sentence structure may be associated with a support query category 81E.

The support query category 81E may be associated with multiple different types of support intent categories and could even be identified as a parent intent category for multiple narrower support query intent categories in an intent category hierarchy. The natural language engine 71 uses the identified support query category 81E along with the <cell phone> concept 81D identified for the noun phrase 81B to identify a cell phone support query intent category 81F for query 70F.

One system that conducts the linguistic analysis described in FIGS. 5A and 5B is the Inquiria Matching Language described in co-pending patent application Serial No. 10/820,341, filed April 7, 2004, entitled: AN IMPROVED ONTOLOGY FOR USE WITH A SYSTEM, METHOD, AND COMPUTER READABLE MEDIUM FOR RETRIEVING INFORMATION AND RESPONSE TO A QUERY, which has already been incorporated by reference in its entirety. Of course, other types of natural language systems could also be used.

FIG. 6 explains further how the intent management tool 67 in FIG. 4 can be used to update intent categories. In operation 80, previous queries are logged for some time period. For example, all of the queries for the past week. In operation 82, the intent management tool compares the logged queries with existing intent categories. Any intent categories matching more than a first threshold number of logged queries may be identified in operation 83A. Matching logged queries with existing intent categories can be performed in a similar manner as described above in FIGS. 5A and 5B. The intent responses for any identified intent categories in operation 83A may then be posted on an enterprise webpage in operation 83B. For example, 20% of the logged queries may have been associated with “contract expiration support” questions. If the threshold for adding an intent response to the enterprise web page is 15%, then a link to information relating to “contract expiration support” may be posted on the enterprise home web page in operation 83B.

Optionally, the intent management tool may in operation 84A identify information currently displayed or listed on the enterprise webpage that have an associated intent category that does not match a second threshold number of logged queries. Information associated with intent categories below this second threshold may be removed from the enterprise
5 webpage in operation 84B. For example, the enterprise home web page may currently display a link to an interest free checking promotion. If the number of logged queries matching an “interest free checking” intent category are below the second lower threshold, such below 1% of all logged queries, the “interest free checking” link or information can be identified by the intent management tool 67 and then either manually or automatically
10 removed from the enterprise web page.

This provides a valuable system for promoting different services or products to users. For example, as described above, the intent management tool 67 can be used to determine that the “interest free checking” promotion is of little interest to customers. Alternatively, the same intent management tool 67 can determine that a “home refinance” promotion associated
15 with a “home refinance” intent category has a substantially larger number of matching queries. Accordingly, a website administrator can quickly replace the interest free checking promotion with the home refinance promotion on the enterprise home web page.

In operation 85A, the software executing the intent management tool 67 may automatically identify frequently queried intent categories that have no associated intent
20 response. For example, the intent management tool 67 may identify intent categories with no currently associated intent response that match a third threshold number of logged queries. In operation 85B, the intent management tool 67 asks the administrator to identify an intent response for any identified intent categories. The intent responses can be information such as a web page and/or links to information on a web page that is responsive to the identified

intent category. The intent responses input by the administrator are then assigned to the associated intent categories by the intent management tool 67 in operation 85C.

In yet another operation 86A, the software operating the intent management tool 67 may identify related queries with no associated intent categories. For example, a group of queries may be identified that are all related with a same financial service promotion but that currently have no assigned intent category. The largest number of related queries with no associated intent category may be identified first, and then lower numbers of related queries listed, etc. Alternatively, the intent management tool 67 can be configured to only list related queries over some predetermined threshold number.

The intent management tool in operation 86B asks the administrator to identify an intent category for the group of identified queries. Alternatively, the common information identified in the group of queries may be used as the intent category. In operation 86C, the intent management tool 67 then asks the user to identify an intent response for the identified intent category.

FIG. 7 shows one way the intent management tool 67 can be used to update existing intent categories. In operation 90, queries are logged for some period of time in the same manner described above in FIG. 6. In operation 92, the intent management tool 67 identifies logged queries that do not match any current intent categories. One or more new intent categories are then created for the non-matching queries in operation 94. The new intent categories are either manually generated by the administrator or automatically generated by a natural language engine 71 as described above in FIGS. 5A and 5B.

The new intent categories are then run against the logged queries in operation 96. This operation may be iterative. For example, the number of matching queries is identified in operation 98. If the number of queries matching the new intent category is below some threshold, such as a predetermined percentage of the logged queries, the new intent category

may be refined in operation 97 and then compared again with the logged queries in operation 96. For example, the administrator may modify certain words in the intent category that may cause more matches with the logged queries. When the new intent category matches more than some threshold number of logged queries in operation 98, IML is created in operation 99
5 that automatically matches the same logged queries with the new intent category.

In another embodiment, the new intent category may also be applied to other query logs to verify accuracy. For example, a query log from another time period may be applied to the newly created intent category. The operation described above for generating new intent categories can also be used when generating the initial intent categories for an enterprise.

10 It should also be understood that an industry expert may be used to review the logged queries and then manually generate useful intent categories based on the results from the intent management tool 67. The fact that the intent categories are “useful” is worth noting. Some clustering algorithms may generate information categories that, for example, may be too broad to really provide useful information. For example, as described above, a clustering
15 algorithm may identify queries all related to “email”. However, providing and supporting a general email intent category may be of little relevance when trying to provide responses to queries directed to an online financial institution.

The industry expert can first derive pertinent intent categories and then refine the derived intent categories to optimize the number of queries matches. This ensures that intent
20 categories are useful and are relevant to the actual queries submitted by users. The optimized intent categories are then used by the search engine to identify query meaning. Alternatively, all or part of the intent category generation can be automated using the intent discovery tool 67 as described above in FIGS. 5A and 5B and as described in further detail below.

The intent discovery tool 67 also allows the web site administrator to identify queries
25 that do not correspond with current intent categories. For example, users may start using new

terminology in queries referring to a new service or product. The intent discovery tool 67 can identify these queries that do not match existing intent categories and then either modify an existing related intent category or create a new intent category that matches the identified queries.

5

Intent Hierarchy

FIG. 8 shows how hierarchies can be associated with intent categories. In this example, a group of queries 100 are all associated with a retirement plan research intent category 110. Either manually or through the intent management tool 67 in FIG. 4, a Roth
10 intent category 102 is derived for a first group of queries 100A, a Regular IRA intent category 104 is created for a second group of queries 100B, and a 401K intent category is derived for a third set of queries 100C.

Again either manually by an industry expert, or automatically with the management tool 67, an intent hierarchy 126 is derived for the intent categories 102-110. For example, a
15 parent “IRA” intent category 108 is derived for intent categories 102 and 104. In addition, a parent “Retirement Plan Research” intent category 110 is derived for intent categories 108 and 106.

This intent hierarchy 126 can be derived in a variety of different ways, but in one example could use clustering analysis as described in more detail below in FIG. 15. A
20 hierarchy tag can then be assigned to the intent categories to control what responses are automatically presented to a user.

To explain further, FIG. 9 shows how the intent hierarchy 126 is used in combination with an identified intent category 102. The intent based search engine may receive a query 120 that matches the Roth intent category 102 previously described in FIG. 8. Accordingly,

the search engine displays an intent response 128A associated with the identified “Roth” intent category 102.

However, the intent category 102 can also be assigned a tag 124 that directs the search engine to display responses for any parents of the “Roth” intent category 102. Accordingly, by selecting tag 124, the search engine refers to the intent category hierarchy 126 to identify any parents of “Roth” intent category 102. In this example, the “IRA” intent category 108 and the “Retirement Plan Research” intent category 110 are identified as parents. Accordingly, intent responses 128B and 128C associated with intent categories 108 and 110, respectively, are also displayed in response to query 120.

Notice that in this example, the intent response 128C associated with parent intent category 110 includes a promotional advertisement for opening a 401K account. Since the “Roth” intent category 102 and the “401K” intent category 106 both have a common parent 110, the enterprise can use tag 124 to promote services, products, or present other information to users that is related to a common broader subject matter than what is actually contained in query 120.

This intent hierarchy feature provides a powerful tool for providing relevant information responsive to a query. For example, a user may not necessarily know they are seeking information related to a 401K account. However, the user is aware of IRA accounts. Intent hierarchy tag 124 allows the search engine to automatically associate a question related to IRAs with an intent category related to 401K accounts based on the classification of both IRA and 401K accounts under the same “Retirement Plan research” parent intent category 110. Thus, the user may receive some relevant 401K information under the “Retirement Plan research” intent category 110 without ever using the word “401K” in the query 120. This also has promotional and advertising advantages. For example, the enterprise can notify any user sending any type of retirement plan related query of a new 401K promotion.

The intent hierarchy tag 124 can consist of a pointer to an associated intent hierarchy 126 as shown in FIG. 9. The intent hierarchy tag 124 can also be used to direct the search engine to display intent responses associated with child intent categories or associated with other intent categories not contained in the same hierarchy.

5

Parameterized Intent Categories

FIG. 10 shows another embodiment of the intent based search engine 35 that allows an administrator or user to associate different parameters with intent categories. The intent management tool 67, for example, may process logged queries 68. In this example, the intent management tool 67 either identifies or creates a “vehicle research” intent category 130 and may then assign different intent responses 134 to the intent category 130 using parameters 135.

The management tool 67 automatically compares the intent category 130 with one or more ontologies 133 and determines that the word “vehicle” 131 in the intent category 130 is associated with the <vehicle> concept 132A in ontology 133. The management tool 67 may then present the user with a drop down menu or, some other type of display, that shows the different concepts or other words or phrases associated with the <vehicle> concept 132A in ontology 133. In this example, the concepts 132A-132E in ontology 133 are displayed as parameters 137A-137E, respectively. The parameters 137A-137E may include pointers to associated intent responses 134A-134E, respectively.

The administrator can select which of the parameters 137A-137E (pointers to intent responses 134) to associate with the intent category 130. In this example, the administrator at least selects minivan parameter 137B. The search engine 35 will then use the assigned parameter 137B to provide additional responses to an associated query. For example, the search engine 35 may later receive a query 139 associated with the vehicle research intent

25

category 130. The search engine 35 identifies the selected parameter 137B assigned to intent category 130 and accordingly displays the intent responses 134B.

In another embodiment, the intent parameters 135 may also cause the search engine to display responses for any associated parent concepts. For example, a query may be associated with a minivan research intent category. A parameter 135 can be assigned to the minivan intent category that causes the search engine to provide responses associated with any of the broader concepts in ontology 133, such as a response related to the vehicle research intent category 130.

The selection of different parameters 135 can similarly be performed by a user. For example, the search engine 35 may initially display the intent category 130 to the user based on a received query. The user can then be provided with the same list of different parameters 137A- 137E associated with the ontology 133. The user then selects which intent responses 134 to display by selecting any combination of parameters 137A-137D.

It is worth noting that the intent category hierarchy described above in FIGS. 8 and 9 and the intent parameters shown in FIG. 10 may be useful in classifying different types of queries. For example, the intent hierarchies in FIGS. 8 and 9 may be better at classifying queries that include more verbs, and the intent parameters in FIG. 10 may be better at classifying queries that include more nouns. For example, questions related to specific types of products may include more nouns while questions related to services or user activities may include more verbs. Of course, these are just examples and either the intent category hierarchy or the intent parameters can be used for any type of query.

Generating New Intent Parameters

Referring still to FIG. 10, the intent management tool 67 can also be used for identifying new intent parameters 140. The intent management tool 67 may identify a large

group of queries all matching intent category 130 but that do not match any of the existing parameters 135 or associated concepts 132 in ontology 133. For example, a group of queries may all be associated with a new minivan model C that is not currently identified in ontology 133.

5 The intent management tool 67 suggests adding a new parameter 137F to parameter list 135 that is associated with the identified minivan model C. Upon selection, parameter 137F is add to parameter list 135. The intent management tool 67 may also ask the administrator to add any other synonyms associated with the new model C parameter 137F and provide an associated intent response 134F. In addition, the intent management tool 67
10 may update ontology 133 to include a new model C concept 132F underneath the minivan concept 132B.

User Classification

 The intent management tool 67 can also assign different “user” related parameters to
15 intent categories. This allows the intent based search engine to associate particular intent responses or search engine actions with different types of users. For example, it may be desirable to automatically initiate a phone call to any long term user that has sent a query associated with terminating an existing account. In another scenario, it may be desirable for the search engine to track the number of times particular users send queries associated with
20 particular intent categories. The search engine can then send intent responses based on the tracked frequency.

 Referring to FIG. 11, any of these different user associated parameters are assigned to particular intent categories by the administrator using the intent management tool 67. The intent based search engine 35 may then receive a query in operation 150. The search engine

identifies an intent category for the query in operation 152 and identifies any user parameters that may be associated with the identified intent category in operation 154.

The search engine in operation 156 conducts any user operation according to the identified user parameters. For example, the user parameter may direct the search engine in operation 158 to track the user query frequency and then classify the user according to the identified frequency. This could be used for providing special promotional materials to high frequency users. Accordingly, the user parameter may direct the search engine in operation 159 to display certain intent responses to the user according to the user classification. The user classifications can also be based on factors unrelated to user query frequency. For example, the user classifications may be based on how long the user has been signed up on the enterprise website; priority membership status, such as a platinum membership, geographic region, age, or any other user demographic.

Intent Discovery

Clustering algorithms are used for statistically associating together different information. A set of features are input into the clustering algorithm which then groups together different information according to the features. These types of conventional clustering algorithms are known to those skilled in the art, and are accordingly not described in further detail.

The present intent discovery scheme may provide different concepts to the clustering algorithm as features that then allow the clustering algorithm to more effectively cluster together related queries. The features provided to the clustering algorithm can be any combination of words, stems, tokens, phrases, concepts, intent categories, etc. FIG. 13 describes this intent discovery scheme in more detail. A set of queries 175 may be input to a clustering engine 186. As opposed to conventional keyword clustering, the clustering engine

186 is given at least a partial set of features 184 associated with the concepts in an enterprise specific ontology 183. For example, the stems, tokens, phrases, and/or concepts in ontology 183 may all be associated with the concept “family vehicle”.

The clustering engine 186 analyzes the queries 175 with respect to the ontology based features 184. Some higher order concepts, such as the concept “family vehicle” may get a larger weight when the queries 175 are clustered than lower order concepts, such as “vehicle models”. The clustering engine 186 outputs names 188 for the identified clusters that, for example, may comprise a string of the most common terms and highest order concepts in the clusters.

Then either through a manual or automated process, IML expressions 190 are created that match the queries in the identified clusters with a particular intent category. The intent categories may use some or all of the terms from the cluster names. For example, the string of most common terms 192 contained in queries 182 may be used in the IML expression 190 to identify station wagon queries 182. Other concepts in ontology 183 can also be used in the IML expression 192 to help classify the station wagon queries 182.

Referring to FIG. 13, the above clustering scheme can also be used to further improve or automate intent classification. For example, the intent management tool 67 described in FIG. 4 may be used in operation 190 to identify any of the logged queries that do not match any of the existing intent categories. In operation 192, the identified queries are submitted to the clustering engine 186 in FIG. 12. In operation 194, features from one or more of the ontologies 183 in FIG. 12 are also fed into the clustering engine 186. The intent management tool 67 receives the names identified by the clustering engine in operation 196 and uses the cluster names and the identified clustered queries to generate new intent categories in operation 198.

Referring to FIG. 14, the intent discovery scheme can also be used to create intent hierarchies. For example, intent category 200 for “family vehicles” and intent subcategory 201 for “minivans” have already been created. However, the intent discovery scheme described above may have discovered three new intent categories 202A-202C.

5 The intent management tool 67 may compare the queries matching multiple intent categories to determine where the new intent categories 202A-202C should be located in the hierarchy. For example, the intent management tool 67 may discover that all of the queries matching new intent category 202C are a subset of the queries matching existing parent intent category 200. Further, the intent management tool 67 may also determine that the queries
10 matching new intent category 202C do not, or rarely, overlap with the queries matching “minivan” intent category 201. Accordingly, the intent management tool 67 locates new intent category 202C as a direct child of intent category 200.

 It may be determined that the queries matching the other new intent categories 202A and 202B are a subset of the queries matching existing intent category 201. Accordingly,
15 new intent categories 202A and 202B are assigned as direct descendants of intent category 201. The intent management tool 67 may also identify new parameters for an existing intent category as described above in FIG. 10.

Response Parameters

20 FIG. 15 shows another type of parameter that can be assigned to different intent responses. An intent response 220 may comprise a template associated with a particular general category of questions. For example, the intent response 220 may be associated with an intent category related to buying and viewing a vehicle. Instead of creating a separate intent response for every specific model of vehicle that a user may ask about, the intent

response 220 may include parameters 222 and 224 that are associated with specific information elements within the query.

For example, response parameter 222A may be associated with price information 228A for a particular minivan model and response parameter 222B may be associated with price information 228C for a particular station wagon model. Similarly, response parameter 224A may be associated with image information 228B for the minivan and response parameter 224B may be associated with image information 228D for the station wagon.

The intent based search engine 35 receives the query 230 and conducts the linguistic analysis described above to determine an associated intent category. The identified intent category is associated with intent response 220. The search engine 35 then compares elements in the query 230 with the response parameters 222 and 224 to determine what additional response elements 228 to insert into intent response 220.

In this example, the search engine matches the <minivan> concept parameters 222A and 224A in intent response 220 with the word minivan in query 230. Accordingly, the response elements 228A and 228B in table 226 are displayed with the intent response 220 on user interface 232. The response parameters allow an almost identical intent response 220 to be generated for all of the queries within a particular intent category and then automatically customize the intent response 220 for different query elements.

The system described above can use dedicated processor systems, micro controllers, programmable logic devices, or microprocessors that perform some or all of the operations. Some of the operations described above may be implemented in software and other operations may be implemented in hardware.

For the sake of convenience, the operations are described as various interconnected functional blocks or distinct software modules. This is not necessary, however, and there may be cases where these functional blocks or modules are equivalently aggregated into a

single logic device, program or operation with unclear boundaries. In any event, the functional blocks and software modules or features of the flexible interface can be implemented by themselves, or in combination with other operations in either hardware or software.

- 5 Having described and illustrated the principles of the invention in a preferred embodiment thereof, it should be apparent that the invention may be modified in arrangement and detail without departing from such principles. I/we claim all modifications and variation coming within the spirit and scope of the following claims.

CLAIMS

1. A method for classifying information requests, comprising:

receiving a group of text questions logged for one or more servers for a period of time;

5 identifying groups of the logged questions that use a variety of different linguistic expressions to express similar information requests;

assigning intent categories applicable to each group of questions requesting similar information; and

configuring intent responses that each provide applicable information to all of the

10 questions assigned to the same intent categories.

2. The method according to claim 1 including:

conducting spelling analysis to replace unidentifiable words in the questions;

conducting punctuation analysis to correct punctuation in the questions;

15 conducting sentence analysis to identify sentence structures and sentence elements in the questions;

conducting a stem analysis to add and identify other forms of words to the questions;

conducting a concept analysis using one or more ontologies to identify concepts related to the questions; and

20 conducting a linguistic analysis using the identified concepts, sentence structures, sentence elements, and other added and identified forms of words to identify intent categories for the questions.

25 3. The method according to claim 1 including:

identifying the intent categories associated with the largest groups of questions; and
posting pre-query information or links on a web page that provide intent responses for
the identified intent categories.

- 5 4. The method according to claim 3 including:
receiving a new group of questions logged for the one or more servers for another
period of time;
using natural language analysis to identify new groups of questions not associated
with existing intent categories; and
10 assigning new intent categories for the identified new groups of questions; and
configuring new intent responses that provide information applicable to all of the new
questions in the same new intent categories

5. The method according to claim 4 including:
15 identifying any of the new intent categories associated with at least a pre-query
threshold number of new questions; and
posting the new intent responses for the identified new intent categories on a web
page.

- 20 6. The method according to claim 5 including:
identifying any existing intent categories that are not associated with at least a lower
threshold number of new questions; and
removing pre-query intent responses for any of the identified existing intent
categories from the web page.

25

7. The method according to claim 7 including configuring at least some of the intent categories into an intent hierarchy that causes the intent responses associated with the intent categories to be displayed according to the configured intent hierarchy.

5 8. The method according to claim 1 including assigning parameters to the intent categories that control how the associated intent responses are displayed.

9. The method according to claim 8 including:

identifying ontologies associated with the intent categories;

10 presenting concepts in the identified ontologies for selection by a user;

assigning any of the selected concepts as ontology parameters for the associated intent categories; and

displaying intent responses for the ontology parameters assigned to the intent categories.

15 10. The method according to claim 8 including associating user parameters with the intent categories that cause a search engine to display different information associated with a user submitting the question or associated with a query history associated with the user submitting the question.

20 11. The method according to claim 1 including assigning response parameters to the intent responses that cause a search engine to display both the intent responses associated with the intent categories and also display any additional information associated with the response parameters assigned to the intent responses.

25

12. A method for discovering query intent categories, comprising:

logging queries received by an enterprise information system;

using a clustering engine to identify clusters of the logged queries;

generating names for the clusters of logged queries;

5 using the generated names to create intent categories pertinent to the queries in the same clusters; and

using a linguistic matching language to match the queries in the same clusters with the intent categories.

10 13. The method according to claim 13 including using ontology elements associated with the enterprise information system as features for the clustering engine and then using at least some of the cluster names generated by the clustering engine to create intent categories.

14. The method according to claim 13 including:

15 identifying queries that do not match any of the created intent categories;

submitting the identified queries to the clustering engine;

generating new cluster names for the clusters of non-matching queries; and

using the new cluster names to generate additional intent categories.

20 15. The method according to claim 14 including:

automatically generating an intent hierarchy by identifying a new intent category

having an associated set of queries that comprise a sub-set of queries for an existing intent

category and identifying the new intent category as a child of the existing intent category; and

identifying a new intent category having a subgroup of associated queries that comprise all of the queries matching an existing intent category and identifying the new intent category as a parent of the existing intent category.

5 16. A search engine, comprising:

a processor configured to receive queries and then conduct a linguistic analysis that identifies different concepts and linguistic formations in the queries;

the processor further configured to identify the received queries having similar information requests according to the identified concepts and linguistic formations and

10 classify the queries with similar information requests under similar intent categories; and

the processor further configured to provide common pertinent information responses to all of the similar information requests classified under the same intent categories.

17. The search engine according to claim 16 wherein the processor is configured to

15 operate an Intelligent Matching Language (IML) engine and use one or more ontologies to identify the concepts and linguistic formations in the queries.

18. The search engine according to claim 17 including a memory that stores

preconfigured intent responses for only a portion of most frequently queried intent categories

20 and returns the preconfigured intent responses for the intent categories associated with the received queries.

19. The search engine associated with claim 18 wherein the processor is configured to

conduct searches using concepts identified using the IML engine and the ontologies when no

intent categories can be identified for the received queries and is further configured to conduct keyword searches when no concepts can be identified in the queries.

20. The search engine according to claim 16 wherein the intent categories are arranged in
5 intent hierarchies and the processor is configured to provide intent responses corresponding to locations of the intent categories in the intent hierarchies.

21. The search engine according to claim 16 wherein the intent categories include
different parameters associated with different intent responses.

10

22. The search engine according to claim 21 wherein the parameters are associated with
ontology concepts and the processor displays intent responses corresponding with the
ontology concepts.

15 23. The search engine according to claim 21 wherein the parameters are associated with types of user classifications or types of user operations performed by users submitting the queries.

24. The search engine according to claim 21 wherein the intent responses include
20 response parameters associated with additional intent responses and the processor is configured to identify elements in the queries associated with the response parameters and display the associated additional intent responses.




FIG.1

2/16




FIG.2




FIG.3A

4/16





FIG.3B

5/16



6/16




FIG.5A

7/16

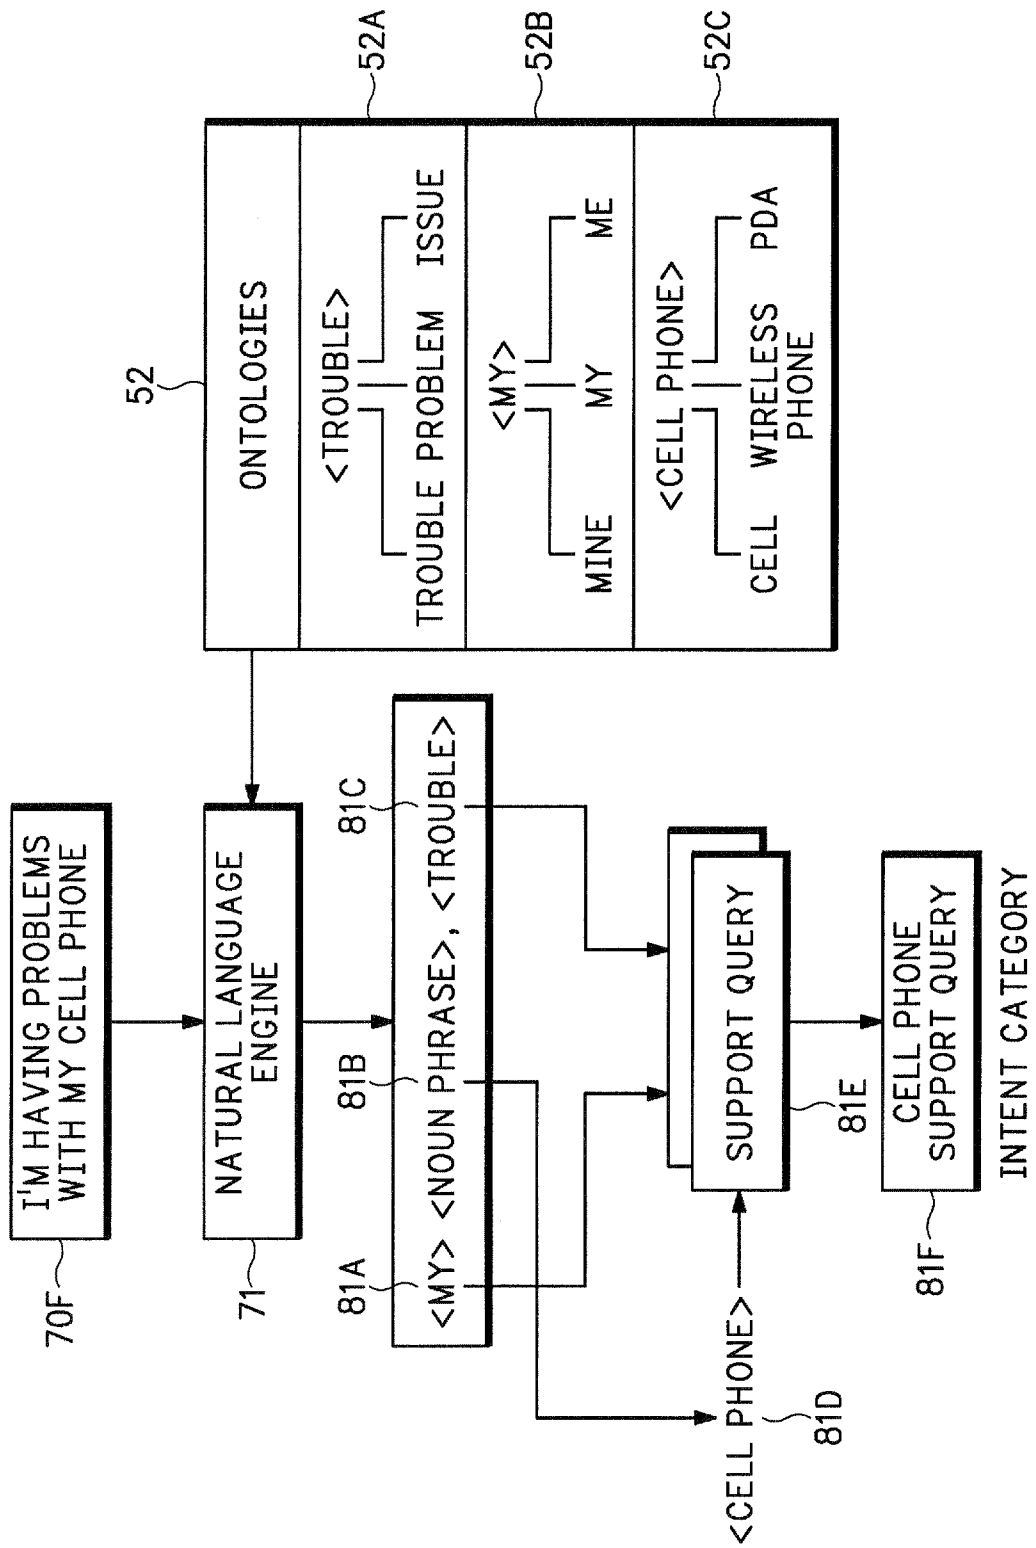


FIG.5B

8/16

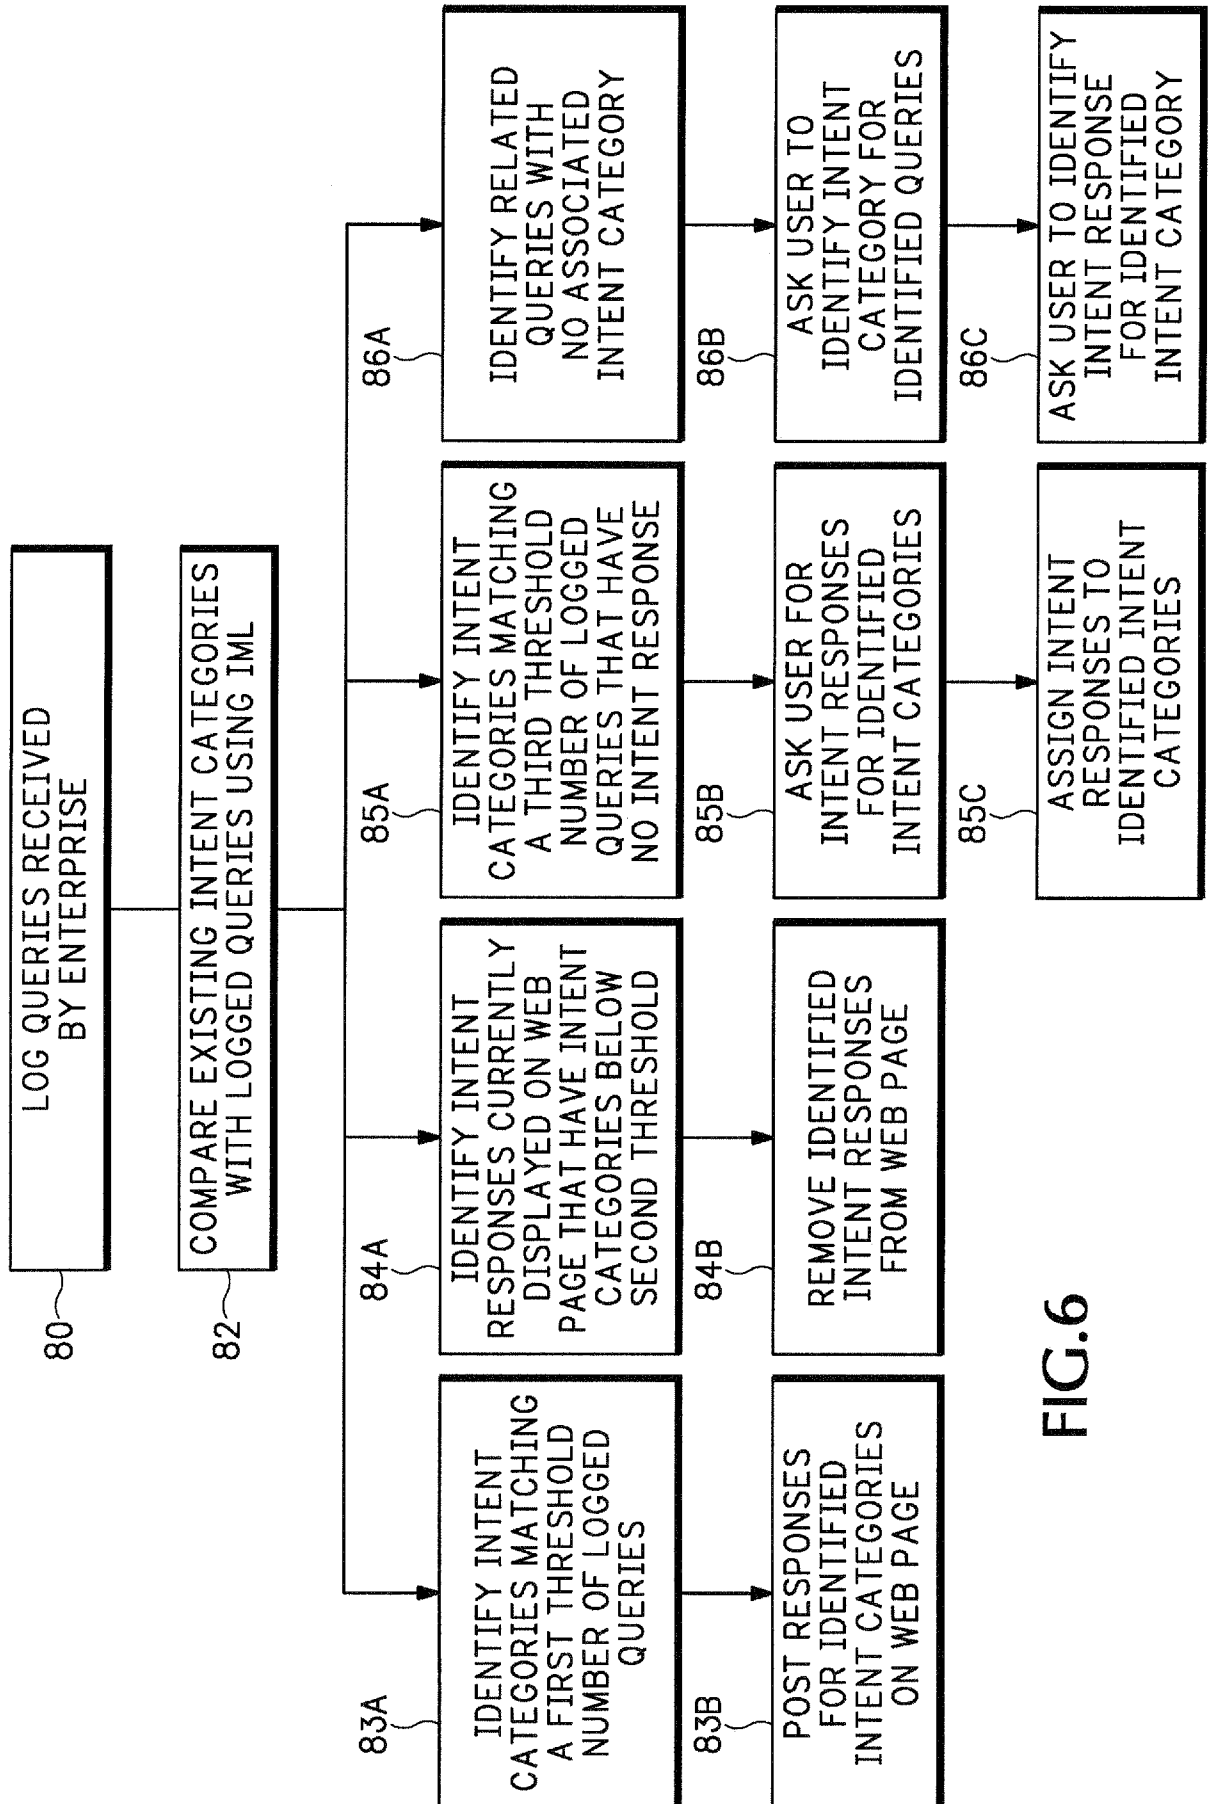


FIG.6

9/16

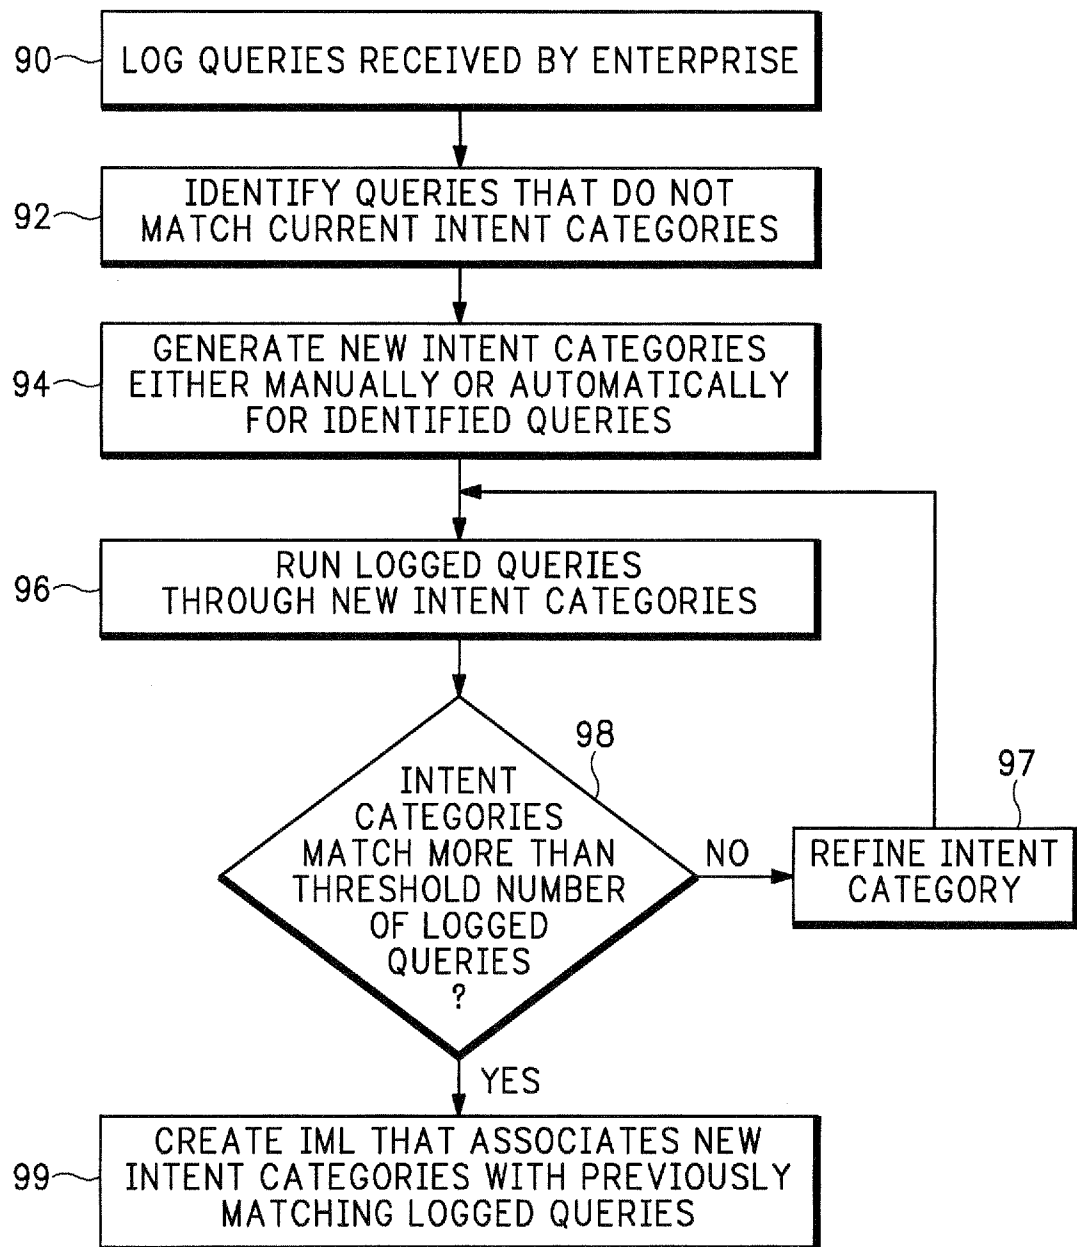


FIG.7

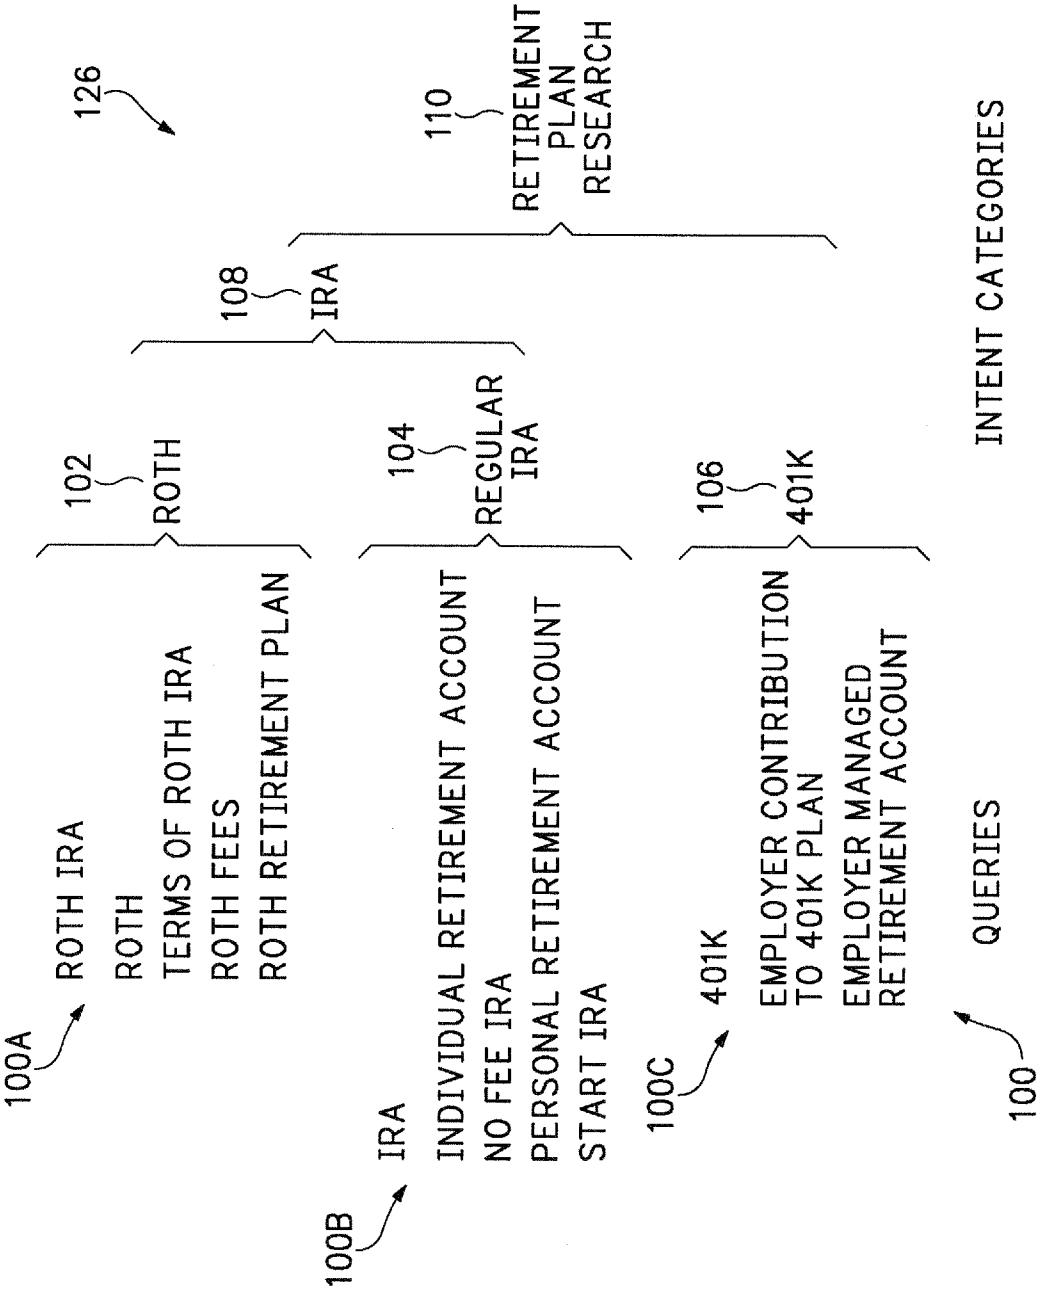


FIG.8




FIG. 9

12/16




FIG. 10

13/16

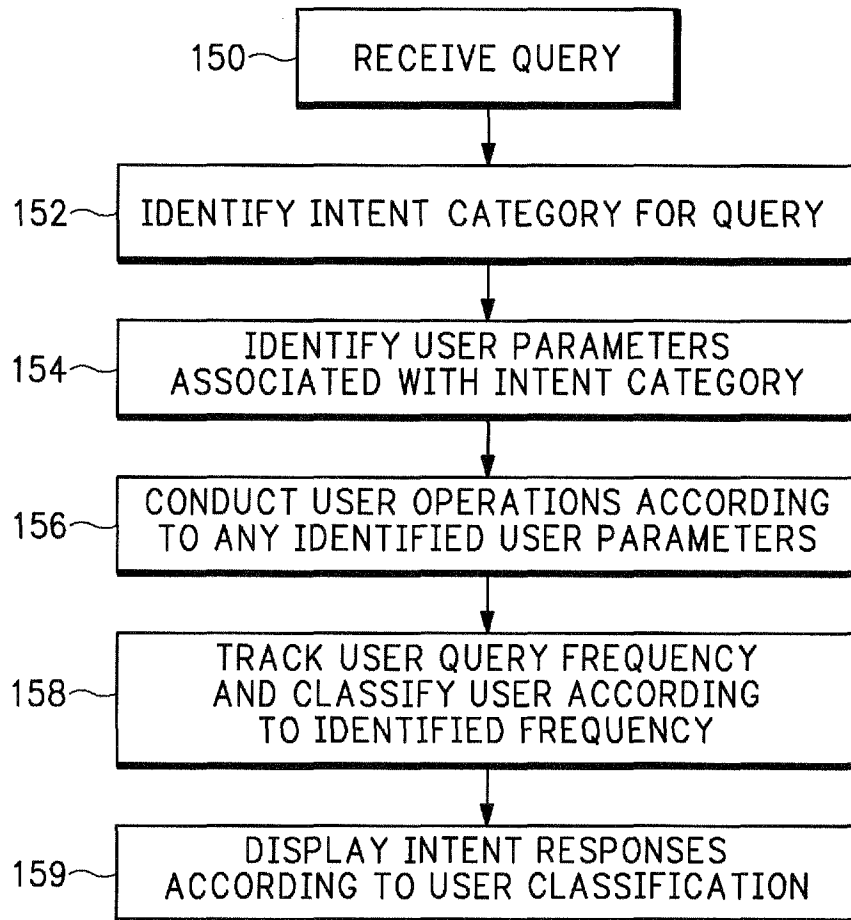


FIG.11

14/16




FIG.12

15/16

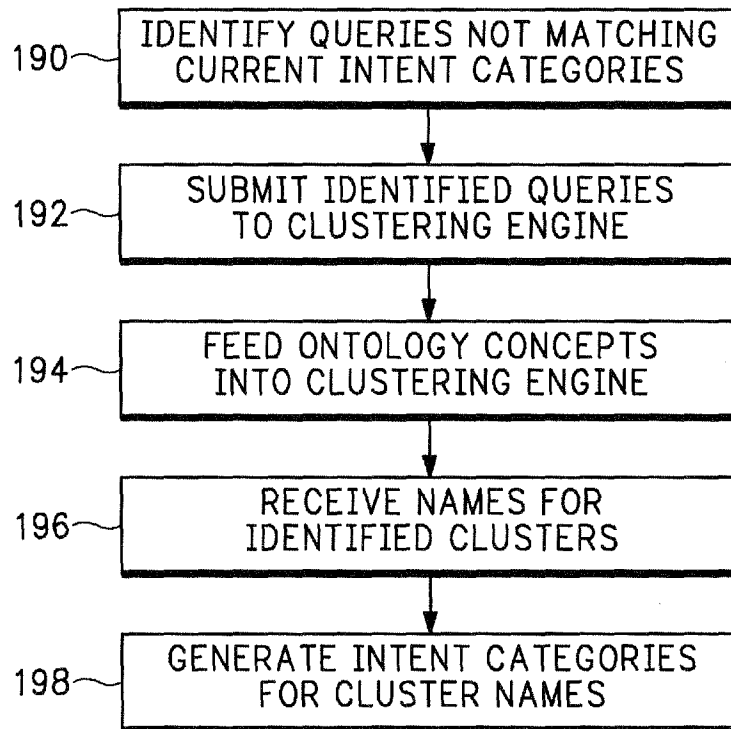


FIG.13

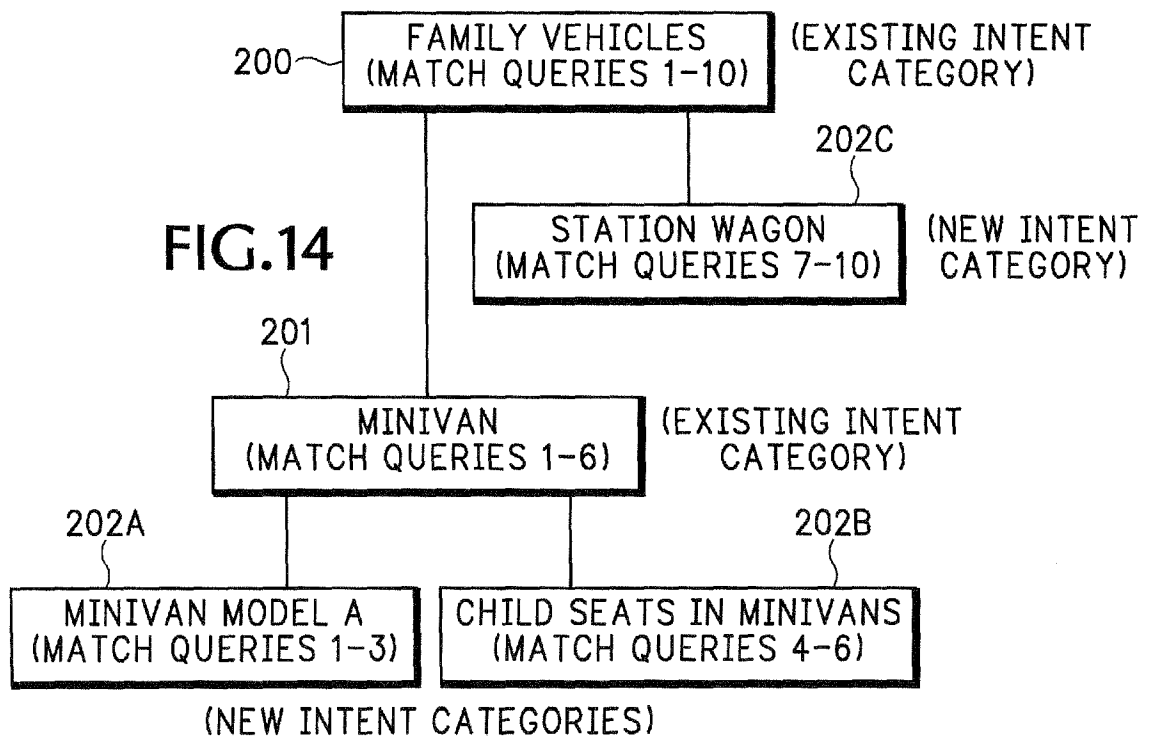


FIG.14

16/16




FIG. 15