




- (51) **International Patent Classification:**
G06N 5/04 (2006.01) *G06F 16/33* (2019.01)
G06F 17/27 (2006.01)
- (21) **International Application Number:**
PCT/US2018/055860
- (22) **International Filing Date:**
15 October 2018 (15.10.2018)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
15/785,354 16 October 2017 (16.10.2017) US
15/785,347 16 October 2017 (16.10.2017) US
- (63) **Related by continuation (CON) or continuation-in-part (CIP) to earlier applications:**
- (71) **Applicant:** NUANCE COMMUNICATIONS, INC.
[US/US]; One Wayside Road, Burlington, MA 01803 (US).
- (72) **Inventors:** YEH, Peter; 159 Newbury Street, Milpitas, CA 95035 (US). STORY, Ezra; 8061 26th Ave Nw, Seattle, WA 98117 (US). JAIN, Prateek; 77 N Almaden Ave, Apt 1814, San Jose, CA 95110-2795 (US).
- (74) **Agent:** SOLOMON, Mark, B. et al.; Hamilton, Brook, Smith & Reynolds, P.C., 530 Virginia Road, P.O. Box 9133, Concord, MA 01742-9133 (US).
- (81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,

(54) **Title:** SYSTEM AND METHOD FOR CONTEXTUAL REASONING



(57) **Abstract:** Current virtual systems, however, are not aware of and do not consider many pieces of contextual, or implicit, information in responding to user requests. Examples of contextual information can include, in a vehicle context, accessibility needs for a parking request, or weather based information for covered parking, or parking closer to a requested destination venue. Such contextual information can be provided by a user settings file, or systems or sensors of a car, such as systems indicating a windshield wiper is running or headlights are activated, and the like. Determining contextual information can be applied to multiple application areas such as providing directions to the most appropriate destination given a user's preferences and sensed information from the vehicle. For example, the system can find covered parking in rain/snow conditions, as opposed to uncovered parking.

CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,
KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, 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, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

SYSTEM AND METHOD FOR CONTEXTUAL REASONING

RELATED APPLICATIONS

This application is a continuation of U.S. Application No. 15/785,347, filed October 16, 2017 and U.S. Application No. 15/785,354, filed October 16, 2017. The entire teachings of the above applications are incorporated herein by reference.

BACKGROUND

[0001] Natural language understanding (NLU) systems receive user speech and translate the speech directly into a query. Often, NLU systems are configured to operate on smartphones. These NLU systems can also direct such a query to a search engine on the smartphone or accessible via a wireless network connection to perform Internet searches based on the content of the query.

SUMMARY

[0002] An embodiment of the present invention, a method, and corresponding system and non-transitory computer readable medium, includes determining, based on a received query and contextual information, candidate reasoners to respond to a portion of the received query. A “reasoner” or “candidate reasoner” is defined herein as a module that translates information from a sensor, information within user settings, information of preferences of the user learned from prior interactions with the user or information from any source related to the user or to the received query, into additional fields or values for a query or the revision of existing fields or values. The method further includes generating, at each candidate reasoner determined, additional query fields based on the contextual information and a rule of a rule database. The method further includes merging the additional or revised query fields or values for each candidate reasoner based on a confidence score associated with a result (or output) of each corresponding candidate reasoner. The confidence score can be based on applicability of the contextual information to the received query. The method further includes providing an enhanced query having the additional or revised query fields.

[0003] In an embodiment, the method and corresponding system and non-transitory computer readable medium further includes determining whether the additional or revised query fields or values should be input to a candidate reasoner to generate further additional or revised query fields or values.

[0004] In an embodiment, the method and corresponding system and non-transitory computer readable medium includes generating the received query based on a natural language understanding request.

[0005] In an embodiment of the method and corresponding system and non-transitory computer readable medium, can be configured to consider a plurality of candidate reasoners including at least one of a spatial reasoning module configured to perform reasoning based on indications of locations of information in the query, a temporal reasoning module configured to perform reasoning based on timing of information in the query, contextual reasoning module configured to perform reasoning based on contextual information, preference reasoning module configured to perform reasoning based on preferences of the users, and custom reasoning modules configured to perform reasoning based on domain specific information present in the context or received query such as reasoning about what cuisines to consider if the requested or preferred cuisine is not available. The candidate reasoner is configured to allow a variety of different reasoners, using a plug-in architecture, that can be added to or removed from the system. The plug-in architecture, in an embodiment, wraps the inputs and outputs so that each module can communicate with the candidate reasoner with similar formats.

[0006] In an embodiment, the method and corresponding system and non-transitory computer readable medium includes gathering contextual information from sensors of a system receiving the query. Examples of sensors include temperature sensors, weather sensors, vehicle systems sensors (e.g., sensors determining whether windshield wiper or headlight systems are active), and the like.

[0007] In an embodiment of the method and corresponding system and non-transitory computer readable medium, merging the additional or revised query fields or values further includes prioritizing the additional or revised query fields or values from a reasoner ranked with a higher priority over additional query fields from a reasoner ranked with a lower priority.

[0008] In an embodiment, the determining, generating, merging, and providing are performed by a reasoning interface layer, and the method and corresponding system and non-transitory computer readable medium further includes providing the enhanced query to an intelligent knowledge layer configured to process the enhanced query including the additional query fields.

[0009] In an embodiment, a system includes a processor and a memory with computer code instructions stored therein. The memory is operatively coupled to said processor such that the computer code instructions configure the processor to implement a reasoning interface layer that is configured to determine, based on a received query and contextual information, candidate reasoners to respond to a portion of the received query, and an intelligent knowledge layer that is configured to determine, based on the output of the reasoning interface layer, the candidate sources to fulfill the query. The system is further configured to generate, at each candidate reasoner determined, additional query fields based on the contextual information and a rule of a rule database. The system is further configured to merge the additional query fields for each candidate reasoner based on a confidence score of each corresponding candidate reasoner. The confidence score is based on applicability of the contextual information to the received query. The system is further configured to provide an enhanced query having the additional query fields.

[0010] In an embodiment, a non-transitory computer-readable medium is configured to store instructions for a reasoning interface layer. The instructions, when loaded and executed by a processor, cause the processor to determine, based on a received query and contextual information, candidate reasoners that are configured to respond to a portion of the received query to select candidate reasoners. The instructions further cause the processor to generate, at each candidate reasoner determined, additional or revised query fields based on the contextual information and a rule of a rule database. The instructions further cause the processor to merge the additional or revised query fields for each candidate reasoner based on a confidence score of each corresponding candidate reasoner. The confidence score is based on applicability of the contextual information to the received query. The weightings further cause the processor to provide an enhanced query having the additional query fields.

[0011] In another embodiment, a method and corresponding system and non-transitory medium include matching terms of a received query to a database of provider capabilities. The method further includes determining, for each portion of the received query, a respective

provider to execute the portion of the received query. The method further includes sending each portion of the received query to its respective reasoner. The method further includes combining results from each respective provider to a returned query result.

[0012] In another embodiment, the method further includes combining the results, based on the results from each provider, filtering a list of results by applying hard constraints to higher priority portions of the received query.

[0013] In another embodiment, determining a respective provider requires receiving results from any of the providers, and, if so, sending appropriate results of the provider to the respective provider as input.

[0014] In an embodiment, a method includes matching portions of a received query to a database of provider capabilities. The method further includes determining, for each portion of the received query, a respective provider to execute the portion of the received query based on the provider capabilities. The method further includes sending each portion of the received query to its respective provider. The method further includes combining results from each respective provider to form a returned query result.

[0015] In an embodiment, combining the results further includes, based on the results from each provider, filtering a list of results by applying hard constraints from portions of the received query, and ranking a list of results based on a priority of soft constraints from portions of the received query.

[0016] In an embodiment, the method includes determining whether a first respective provider requires receiving results from a second provider, and, if so, sending appropriate results of the second provider to the first provider as input.

[0017] In an embodiment, the respective providers include at least one of the following: a third party content provider application programming interface (API), internal database, and a virtual assistant.

[0018] In an embodiment, the method further includes providing a semantic description corresponding to each respective provider. Determining a respective provider to execute the portion of the received query includes determining whether the portion of the received query includes a request for a resource described by the semantic description corresponding to the respective provider.

[0019] In an embodiment, the method further includes constructing a plan that includes an order of invoking each respective provider.

[0020] In an embodiment, combining results from each respective provider to return a query result further includes determining a first result from a first provider, and sending the first result to a second provider as part of the portion of the query to the second provider. In this embodiment, the first result from the first provider is employed to resolve the portion of the query to the second provider.

[0021] In an embodiment, the method includes relaxing a portion of the query if a combination of fields or values of the query produce an incompatible result.

[0022] In an embodiment, the method includes comparing variables and connections between variables of a received query to variables and connections between variables to at least one stitching template to find a best matching stitching template. The method further includes finding a connection present in the best matching stitching template from a first template variable to a second template variable that is absent in the portion of the received query between a first query variable and a second query variable. The first query variable and second query variable correspond to the first template variable and second template variable. The method further includes copying at least one connection from the stitching template to the portion of the query between the first query variable and second query variable.

[0023] In an embodiment, a system includes a processor, and a memory with computer code instructions stored therein. The memory is operatively coupled to said processor such that the computer code instructions configure the processor to implement a planning module configured to match portions of a received query to a database of provider capabilities, and determine, for each portion of the received query, a respective provider to execute the portion of the received query based on the provider capabilities. The processor is further configured to implement a query execution module configured to send each portion of the received query to its respective provider, and combine results from each respective provider to form a returned query result.

[0024] In an embodiment, a method includes determining, based on a received query and contextual information, candidate reasoners to respond to a received query to select a candidate reasoner. A reasoner or candidate reasoner is a module that translates information from a sensor, user settings, or other source, into additional or revised fields for a query. The method further includes generating, at each candidate reasoner determined, additional or revised query fields based on the contextual information and a rule of a rule database. The method further includes merging the additional query fields for each candidate reasoner based

on a confidence score or other metric of each corresponding candidate reasoner. The confidence score can be based on applicability of the contextual information to the received query. The method further includes providing an enhanced query having the additional or revised query fields. Such a method enhances natural language understanding by supplementing queries with contextual information.

[0025] In an embodiment, a method includes matching terms of a received query to a database of provider capabilities. The method further includes determining, for each portion of the received query, a respective provider to execute the portion of the received query. The method further includes sending each portion of the received query to its respective provider. The method further includes combining results from each respective provider to a returned query result.

BRIEF DESCRIPTION OF THE DRAWINGS

[0026] The foregoing will be apparent from the following more particular description of example embodiments, as illustrated in the accompanying drawings in which like reference characters refer to the same parts throughout the different views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating embodiments.

[0027] Fig. 1 is a block diagram illustrating an example embodiment of the present invention.

[0028] Fig. 2 is a block diagram illustrating an example embodiment of a reasoning interface layer.

[0029] Fig. 3 is a block diagram illustrating another embodiment of the reasoning interface layer.

[0030] Fig. 4 is a flow diagram illustrating an example embodiment of a method employed by the present invention.

[0031] Fig. 5 is a diagram illustrating an example embodiment of an intelligent knowledge layer/core engine.

[0032] Fig. 6 is a diagram illustrating an example embodiment of the intelligent knowledge layer.

[0033] Fig. 7A is a diagram illustrating an example embodiment of a query representing a voice request.

[0034] Fig. 7B is a diagram illustrating an example embodiment of identifying candidate plugins based on the query illustrated in Fig. 7A.

[0035] Fig. 7C is a diagram illustrating an example embodiment of constructing a plan to resolve the query by resolving a location of a city, for example.

[0036] Fig. 7D is a diagram illustrating a subsequent iteration to an iteration of Fig. 7C.

[0037] Fig. 7E is a diagram illustrating an embodiment of stitching variables associated with embodiments of Figs. 7A-D.

[0038] Fig. 7F is a diagram illustrating a third iteration of constructing the plan.

[0039] Fig. 7G is a diagram illustrating an example embodiment of stitching variables.

[0040] Fig. 8 is a diagram illustrating an example embodiment of a method of implementing the intelligent knowledge layer employed by the present invention.

[0041] Fig. 9 illustrates a computer network or similar digital processing environment in which embodiments of the present invention may be implemented.

[0042] Fig. 10 is a diagram of an example internal structure of a computer (e.g., client processor/device or server computers) in the computer system of Fig. 9.

DETAILED DESCRIPTION

[0043] A description of example embodiments follows.

[0044] When a user requests information from a virtual system, two problems can arise. First, the user may assume that the system is aware of, or can process, implicit or contextual information. The user's assumption is based on a human's ability to perform reasoning based on contextual information. Contextual information can include environment or personal preferences. Current systems, however, are not aware of and do not consider many pieces of contextual, or implicit, information. Examples of contextual information can include, in a vehicle context, accessibility needs for a parking request, or weather based information for covered parking, or parking closer to a requested destination venue. Such contextual information can be provided by a user settings file, or systems or sensors of a car, such as systems indicating a windshield wiper is running or headlights are activated, and the like.

[0045] Fig. 1 is a block diagram 100 illustrating an example embodiment of the present invention. With reference to Fig. 1, a system 101 receives a voice request 104 for processing. A person of ordinary skill in the art can recognize that the voice request 104, in other embodiments, can be a text-based request, or a formatted query, but for purposes of

simplicity, the description of various example embodiments herein refers to any such user request as a voice request. A person of ordinary skill in the art can understand that the voice request 104 can be processed by an automated speech recognition (ASR) or natural language understanding (NLU) system.

[0046] A reasoning interface layer 102 receives the voice request 104. The reasoning interface layer is configured to perform contextual reasoning to infer implicit constraints, needs, etc. to deliver personalized results to the user issuing the voice request 104. The reasoning interface layer interfaces with multiple reasoning modules, such as a contextual reasoning module 106 and a spatial reasoning module 108 to produce an updated query 110. The reasoning interface layer 102 is described in further detail in reference to Figs. 2-4 below.

[0047] Once such contextual information is determined, the virtual system has to process the information and fulfill that request. The virtual system has to determine which resource to access for a query with multiple parts. The system can have access to competing providers (e.g., sources of information or content) from which to select, or multiple providers that provide outputs that need to be combined together to create the final answer.

[0048] An intelligent knowledge layer 112 receives the updated query 110. A person of ordinary skill in the art can recognize that the intelligent knowledge layer 112 can also operate on a query received directly from the user, or from another source other than the reasoning interface layer 102. However, the reasoning interface layer 102 and intelligent knowledge layer 112 working in combination provide a greater benefit by updating the query and enhancing the result in the intelligent knowledge layer 112.

[0049] The intelligent knowledge layer 112 considers explicit constraints (or hard constraints) in the received query (e.g., updated query 110) to retrieve best results that fulfill the user's request. The intelligent knowledge layer 112 is configured to analyze the updated query 110, determine a strategy for resolving the query using third party sources such as a point of interest (POI) database 116a, parking database 116b, fuel range resource 116c, and restaurants 116d. The intelligent knowledge layer further interfaces with a knowledge repository 118. The intelligent knowledge layer 112 provides an enhanced result 120 to the user. The intelligent knowledge layer 112 is described in further detail in relation to Figs. 5-8 below.

[0050] Solving these problems can be applied to multiple application areas. For example, a voice response system in a vehicle can provide directions to the most appropriate parking garage given a user's preferences and sensed information from the vehicle. For example, finding covered parking in rain/snow conditions is preferred over uncovered parking.

[0051] Another application is with devices using the Internet of Things (IOT) framework. A centralized hub can serve as the gateway between the user and the outside world, where the outside world are services/virtual systems that are accessed through the same hub. For example, a user might request to order a pizza for a Boston Red Sox game. The centralized hub can process the request to determine the time of the Boston Red Sox game, and automatically generate the pizza order to arrive at that time. However, this requires the system to determine when the Red Sox game is, what type of pizza the user likes, where the best restaurant to order from is, whether that restaurant delivers to the user's current address, and how long the restaurant will take to order. After determining this information, it can generate the request to the best restaurant for delivery to the user.

[0052] Applicant's system and method provides advantages by adding contextual reasoning non-query expressed data to an expressed query. While current systems may consider geo-location in contextual reasoning, Applicant's system adds in additional contextual information by using specialized logical reasoners.

[0053] Further, current systems, such as Apple's Siri[®], can be configured to perform a hard wired search to go to a particular provider (e.g., Yelp for restaurants) no matter the query parameters. Each virtual assistant has a dedicated source to use for each type of query. However, depending on the request, it can make more sense to go to a different source particular to the request. Applicant's system employs dynamic selection of sources. Each source is correlated with a declaration of its abilities, and can therefore be matched up better with particular queries. Each declaration states what each source knows about and what each source can provide. Further, each source can be associated with dimensions that it knows about, weights, or confidence scores, of how useful each dimension is, and confidence scores corresponding to each query. For example, when there are two competing sources for a restaurant query, both sources may have information about restaurants and their opening/closing hours, but only one knows more about amenities (e.g., Wi-Fi, handicapped accessible restrooms, type of lighting, bar seating, etc.) at the restaurant. A description for sources captures the known amenities with a high confidence score. Therefore, when user

generates a query about a restaurant with particular amenities, that query can be directed to that source. Another source might have a high confidence score for an availability dimension. Therefore, when a user asks about whether a table is available at a particular time, it is directed to that source. Applicant's system and method therefore has a matching and planning process that compares the information needed to resolve a query against capability descriptions of each available source. Further, a query may include information that is across multiple categories. An example of such a query is "find me a good restaurant near the marina, and directions to parking there." The system needs to first find the location of the marina, then good restaurants (e.g., above a certain rating) near that location, and then parking near the restaurant. This pulls from a database correlating businesses to locations/addresses, another database correlating restaurants that are near a location, and then a third parking database.

[0054] In other words, when asking to park near a destination, the system resolves the first point of interest (POI) before you performing the next search of parking. Other systems may send the unresolved POI to the content source directly, however, sending "restaurant near the marina" to a parking database may be unsuccessful without pre-processing. When the user says "find parking at a good Italian restaurant by the marina," the system needs to resolve the POI "good Italian restaurant by the marina" before sending that to the parking service, because the parking service may not answer the query correctly without a resolved POI.

[0055] The system and method further needs to have a mechanism to perform relaxation to an over-constrained constraint of a query. In other systems, an overly-constrained query can default to an Internet web search of the query, or drop recognized constraints arbitrarily until a result is found. However, each of the constraints can have an importance to the user. That importance can be derived and quantified from the user's preferences or by other contextual cues that the system can infer and assert. For example, the query "find a cheap parking handicapped parking place" may need to determine which is more important – cheap parking, or accessible parking. The system and method can determine the subset of constraints that can be satisfied that brings the most satisfaction to the user. This can be judged by a user-by-user basis. For example, a user that cannot use non-accessible parking under any circumstances is likely to pay whatever is necessary for an accessible parking spot,

despite requesting cheap parking. Therefore, the system prioritizes, in this case, accessible parking over cheap parking, instead of arbitrarily choosing between the two.

[0056] Fig. 2 is a block diagram illustrating an example embodiment of a reasoning interface layer 202. The reasoning interface layer 202 can be a further embodiment of the reasoning interface layer 102 of Fig. 1. With reference to Fig. 2, the reasoning interface layer 202 receives a semantic query 204. The semantic query 204, in this one example, is based on the voice request “find parking near the stadium.” The semantic query can be logically represented as “Parking (x) & Stadium (Y) & Near (X,Y),” where (X) represents the parking location and (Y) represents the stadium location, and “Near(x,y)” resolves a location near (X) and (Y). The reasoning interface layer communicates with a reasoning layer, which can include multiple reasoning modules. In the example of Fig. 2, the reasoning layer includes a contextual reasoning module 206, a spatial reasoning module 208, a temporal reasoning module 214, and a commonsense reasoning module 212. The reasoning interface layer 202 can also receive additional information from sensor application programming interfaces (APIs) 218 of a vehicle 216 or other system. As such, the reasoning interface layer 202 is a flexible framework that allows for a wide range of reasoning techniques to be integrated and accessed via a unified interface. The reasoning interface layer 202 interfaces with a reasoner arbitration which automatically determines which reasoners to employ. Further, the reasoning interface layer 202 can interface with a consistency checker that merges inferences from multiple reasoners into one consistent solution before providing the new query 210.

[0057] In this example, the reasoning interface layer 202 employs the spatial reasoning module 208 and commonsense reasoning module 212. The reasoning engines support frequently occurring reasoning requirements such as spatial reasoning 208, contextual reasoning 206, temporal reasoning 214, and common sense reasoning 212. In this example, to find parking near the stadium, the reasoning interface layer interfaces with the sensor APIs 218 of the vehicle 216 to determine that there is precipitation. From that, the contextual reasoning module 206 determines the new query 210 should include a request for covered parking, and therefore adds “isCovered (x, true)” to the query, indicating that the parking should be covered. Further, the common sense reasoning module 212 determines that distance from the parking to the stadium should be less than a half of a mile, and adds in the Distance (x,y,z) & $z < 0.5$ miles to the new query 210. Therefore, the new query 210 is

enhanced to better serve the user, without the user having to explicitly request these additional query elements.

[0058] Fig. 3 is a block diagram 300 illustrating another embodiment of the reasoning interface layer 302. The reasoning interface layer 302 receives inputs 304 including a query (e.g., a SparQL query), contextual information, and an Artificial Intelligence (AI) Log Object. Then, the reasoning interface layer 302 provides an output 310 of an update query, a provenance, an updated context, and an updated AI Log Object. The reasoning interface layer 302 interfaces with a context mapper 306, reasoner arbitration 322, consistency checker 324, and a meta reasoner 326. The reasoning interface layer 302 further interfaces with a rule repository 318 as well.

[0059] The context mapper 306 converts incoming contextual information from JavaScript Object Notation (JSON) to an internal representation that is utilized by the different reasoners. The context mapper 306 validates the JSON, and if the input JSON is valid, the context mapper, using an open source library, parses the JSON to identify the different keys and their corresponding values. After identifying the keys and values, context mapper 306 converts the keys and values to the internal representation.

[0060] The reasoner arbitration/arbitrator 322 identifies the reasoners that are relevant for the given query and contextual information. A number of arbitration strategies are provided with a reasoning framework and additional arbitration strategies can be added by implementing a Java, or other language's, interface. Some of the arbitration strategies which are available as part of the reasoning framework are Broadcast, Naïve Bayes, and TF-IDF. However, the reasoning framework can further plug-in any arbitration strategy. The plug-and-play of different arbitration strategies is a novel aspect of Applicant's present disclosure, and allows for different systems to easily change their method of arbitration among different reasoners.

[0061] A Broadcast arbitration strategy sends the query and contextual information to all the reasoners plugged into the system.

[0062] A Naïve Bayes arbitration strategy converts the reasoner selection problem into a classification problem and by utilizing a Naïve Bayes classifier for identifying the relevant reasoners. The Naïve Bayes classifier is a supervised training method which receives, as input, a set of features and associated class labels. The definition of a feature depends on the

problem and how a user wants to model them. For example, in a problem where the objective is to identify spam emails, the presence/absence of a word can serve as feature.

[0063] Below is an example embodiment of a rule in XML format that would be stored in the rule repository 318. However, a person of ordinary skill in the art can understand that the rule can be in other formats, including formats other than XML.

```
[0064] <rule>
[0065]   <id> DDISCv4_BIZ_NEED_RESTROOM </id>
[0066]   <gloss>
[0067]     Prioritize coffee shops and fast food restaurants if driver needs
[0068]     to use restroom.
[0069]   </gloss>
[0070]   <reasoner>Business_Reasoner</reasoner>
[0071]   <priority> 4.0 </priority>
[0072]   <active> true </active>
[0073]   <definition>
[0074]     ?PoB rdf:type          nuan:business.generic .
[0075]     ?PoB nuan:fulfills_need ?NEED .
[0076]     FILTER (?NEED = nuan:amenity.toilet)
[0077]   ->
[0078]   Assert:
[0079]     ?PoB nuan:category ?CATEGORY .
[0080]     FILTER SOFT 70
[0081]       (?CATEGORY = nuan:business.coffee_shop ||
[0082]       ?CATEGORY = nuan:business.fast_food)
[0083]   </definition>
[0084] </rule>
```

[0085] With rules such as above, feature construction can be performed as follows: First, the reasoning framework constructs a vocabulary V, which is a set of terms utilized in the rule definition. Then, the reasoning framework constructs a vector for the rules iteratively, marking each term present in V as 0 if the term is absent from the rule. If the term is present in the rule, the term frequency is utilized instead of 0. The reasoner to which the rule is assigned to serves as the class label.

[0086] For example, in case of the rule above, the terms of the rule definition are:

[0087] ?PoB rdf:type nuan:business.generic .

[0088] ?PoB nuan:fulfills_need ?NEED .

[0089] FILTER (?NEED = nuan:amenity.toilet)

[0090] ->

[0091] Assert:

[0092] ?PoB nuan:category ?CATEGORY .

[0093] FILTER SOFT 70

[0094] (?CATEGORY = nuan:business.coffee_shop ||

[0095] ?CATEGORY = nuan:business.fast_food)

[0096] The reasoning framework parses the rule to identify the rule head, which is denoted by the part to the left of/prior to “->” in the rule definition. The rule head is then tokenized which results in a list, such as [?POB, rdf:type, nuan:business.generic,...].

[0097] In the feature vector, the reasoning framework calculates the frequency for each of the terms present in the list. The terms which are present in V, but are not present in the list are marked “1” to account for Laplace Smoothing. The reasoners that the rule belongs to (as specified in the reasoner field - Business_Reasoner) serve as the label for this feature vector.

[0098] To summarize, at the end of the process, there are as many feature vectors as there are rules. Each vector has associated labels, which are the name of the reasoners it is associated with/belongs to.

[0099] The method train the Naïve Bayes method using the feature vectors calculated, as described above. The output is a trained model, which can be utilized on unlabeled data in the next step.

[00100] As described above, the reasoning framework receives a query, which can be a SPARQL query, and contextual information as input. As described herein, the reasoning framework can receive any type of query, but the disclosure herein may refer to SPARQL queries. A person of ordinary skill in the art can recognize that such SPARQL queries can be queries formatted in other ways, including queries represented as first-order logic, etc. The reasoning framework combines the two and converts them into a feature vector using the above described process. The generated vector V' does not have a label or is the unlabeled data. However, the model described above can generate, as output, labels and their

associated probability. The label names generated by the model are the names of the reasoners which are best equipped to process the query and the given contextual information.

[00101] The reasoning framework can employ one of the two strategies to pick the right reasoner. First, the reasoning framework can utilize the top k reasoners by score. Second, a cut off score can be provided to the reasoning framework, and reasoners above the threshold can be utilized for the purposes of a processing a query.

[00102] Term Frequency – Inverse Document Frequency (TF-IDF) is a technique utilized to identify how important a given word is to a document. The TF component indicates how many times a word appears in a given document. The IDF component normalizes the TF component, by calculating the reciprocal of the number of documents in the collection containing the term. For reasoning framework arbitration, TF-IDF arbitration is utilized as follows.

[00103] First, the collection of rules is parsed to remove stop words and identify the term frequency for each rule term (a) in a given rule (b) in the entire rule base. The reasoning framework populates three data structures H1 and H2 to store (a) and (b). H1 is a list of all the rules and a mapping of terms in the rule to their frequencies. H2 consists of all the terms and their frequency. In addition, a dictionary that maps the rule ID to the list of reasoners it belongs to is stored in a data structure D.

[00104] For a new query and contextual information, the method identifies non-stop words, and retains and stores them in a list L . To calculate TF-IDF for each term t in L , there are two components – (i) t 's frequency in every rule and (ii) t 's frequency across the rule base. For each term stored in list L , term frequency – tf in a given rule is calculated by using the structure H1. Term frequency across the rule base - idf is calculated using the structure H2. TF-IDF score (tf-idf) is then calculated using the two components tf and idf.

[00105] Then, the sum of the (tf-idf) for all input query and context terms is utilized to calculate the final tf-idf score for the given query and rule. After the TF-IDF score is calculated using the structure D, the final selection of reasoners can be made in one of the two ways: (1) find the top k rules ranked by score, or (2) identify rules above a specific cut off. Both strategies can be configured in the reasoning framework configuration file.

[00106] The consistency checker 324 employs reasoner mediation to merge the responses returned by the different reasoners. Each reasoner in the reasoning framework has an associated priority with it and this priority is utilized in strategies employed by with the

reasoning framework. A first strategy is a greedy Strategy, which uses the result from the reasoner with highest priority, while other reasoners are ignored. A second strategy is a smart merge strategy, which uses the result from the reasoner with the highest priority is utilized as a starting point for the query to be returned. In other words, the smart merge strategy chooses from a descending order of the priority of reasoners to merge non conflicting components of the result from the other reasoners into the query to be returned. Conflicting components of the results are discarded.

[00107] Fig. 4 is a flow diagram 400 illustrating an example embodiment of a method employed by the present invention. The method begins (402) and determines candidate reasoners to invoke, which can be based on the arbitration strategies disclosed above, including Broadcast, Naïve Bayes, and TF-IDF (404). As described above, multiple candidate reasoners provide flexibility to query enhancement. Additional reasoners can be added with minimal change to the overall structure of the reasoning interface layer. Further, each reasoner can add new types of modular reasoning to the reasoning interface layer, where a single reasoner would have to be re-coded. If no candidate reasoners are found (406), the method ends (418).

[00108] However, upon finding candidate reasoners (406), the method selects the best reasoners to invoke (408) and invokes the reasoner via its reasoner handler/API (410). Then, the method merges conclusions from each invoked reasoner (412) and resolves inconsistencies and conflicts in the merged conclusion (414). This resolution can be based on a predefined strategy such as giving preference to a reasoner with higher confidence. Then, the method determines whether new conclusions were made (416). If so (416), the method begins another iteration. If not, (416), the method ends (418).

[00109] Fig. 5 is a diagram 500 illustrating an example embodiment of an Intelligent Knowledge 512 intelligent knowledge layer/core engine. The intelligent knowledge layer 512 receives an input 510 including a SparQL Query (e.g., from the reasoning interface layer 102 of Fig. 1) and provides an output 520 including a list of results and an updated log object. A person of ordinary skill in the art can recognize that other formats other than the SparQL Query can be used. However, for the purposes of simplicity, SparQL Query/Queries are referenced herein, but can be any other format of query. The intelligent knowledge layer 512 fulfills user requests through semantic routing capabilities and processing of complex to produce high-quality, on-target results. The intelligent knowledge layer 512 provides many

features. First, the intelligent knowledge layer 512 allows for provider arbitration by analyzing each request and determining which provider to invoke to resolve each query. The intelligent knowledge layer 512 further provides knowledge fusion that dynamically chains providers when appropriate and merges their results. A constraint satisfaction and relaxation module of the intelligent knowledge layer 512 further applies implicit and explicit constraints, including logical operators, to perform constraint satisfaction, but further can relax and prioritize constraints as needed to provide constraint relaxation. For example, hard constraints can be applied to high priority portions of the received queries. The hard constraints require that the high priority portions of the received queries are resolved first, followed by lower priority portions of the received query. The constraint satisfaction and relaxation module can further relax constraints by determining portions of the queries that can be ignored in order to return a result. Constraints are relaxed when a result is not possible that satisfies all portion of the query. A data inference module 528 of the intelligent knowledge layer 512 can further infer missing values and attributes based on declarative inference rules. The intelligent knowledge layer 512 can also integrate new content providers using a configurable, declarative framework, using a provider description & plugin manager 532, which includes instructions for interfacing with each provider plugin 534a-n, and also a description of what type of data each provider plugin 534a-n can receive and output.

[00110] A provider arbitration and planning module 522 of the intelligent knowledge layer 512 can analyze the SparQL Query of the input 510 and can determine which providers 534a-n are needed to resolve each part of the query. The provider arbitration and planning module 522 can determine chaining of providers as well. For example, resolving the query “find me handicapped accessible parking by the marina” can require using multiple providers, and chaining the results from each. A first provider can provide a geolocation for the marina, and a second provider can provide a list of handicapped accessible parking facilities near that geolocation. Another example query can be “find me parking near a highly-rated Italian restaurant near the stadium.” Again, a first provider returns a geolocation of a stadium. Then, a second provider returns a four- or five-star Italian restaurant near that geolocation. Then, a third provider searches for parking near the geolocation of the Italian restaurant. By analyzing the query ahead of time using the provider arbitration and planning module 522, these chained requests can be executed accurately.

[00111] The query execution and knowledge fusion module 524 is responsible for executing the plan provided by the provider arbitration and planning module 522 by interfacing with the provider plugins 534a-534n and retrieve the relevant third party content 538a-n. As chained requests are completed, the query execution & knowledge fusion module 524 fuses the information together to determine the result.

[00112] Fig. 6 is a diagram 600 illustrating an example embodiment of the intelligent knowledge layer/core 512. The intelligent knowledge layer 512 receives a SparQL Query 510 and outputs result rows 520, as described in relation to Fig. 5. In relation to Fig. 6, responsive to receiving SparQL Query 510, the intelligent knowledge layer/core 512 generates Queries 640 to be sent to plugins 634a-n. Each plugin then access one or more providers 638a-n to retrieve results 642 to return to the intelligent knowledge layer/core 512. For example, the parking plugin 634a can access Parkopedia 638a, but may arbitrate between several parking providers 638a, depending on the requested information in the queries 640 and the available information advertised by each parking provider 638a. For example, one parking provider may provide information regarding handicapped accessibility, while another provides information regarding amenities at the parking location. Therefore, depending on the query, the parking plugin 634a may be able to better utilize one provider 638a over another provider 638a. Therefore, a query requesting handicapped accessibility should be directed to the parking provider with information about handicapped accessibility. A query requesting amenities at the parking location should be directed to the provider regarding amenities at the parking location. However, a query that requests both information about handicapped accessibility and amenities at the parking location should be arbitrated to determine which information is more important to the user, and direct the query appropriately. In this case, it is likely that handicapped accessibility has a weighting of high importance to the user, and therefore the query is directed to the provider with information about handicapped accessibility.

[00113] Likewise, the location plugin 634b, fuel plugin 634c, and other plugins 634n interface with respective provider(s), such as a HERE geolocation provider 638b, an OPIS fuel provider 638c, and other providers 638n.

[00114] Fig. 7A is a diagram 700 illustrating an example embodiment of query 710 representing a voice request. The voice request, in this example, is “Find disabled parking near city hall in san Francisco.” The query selects “?PID” (Place Identification), “?NAME”

(place name) and “?PRICE” (price of place) from a table given the criteria of the query. The query then needs to solve for the unsolved fields given the defined variables. The query also establishes a soft filter for amenities including an elevator. The soft filter, however, is not a hard requirement of the query, but rather a way to sort results of the other hard requirements. On the other hand, the query establishes a filter for the occupant being a disabled person.

[00115] Fig. 7B is a diagram 720 illustrating an example embodiment of identifying candidate plugins based on the query illustrated in Fig. 7A. Fig. 7B illustrates that to solve for ?PID, the system uses a Parking plugin, for ?POI the system using POISearch, and for ?CITY the system uses a Location/geolocation plugin.

[00116] Fig. 7C is a diagram 750 illustrating an example embodiment of constructing a plan to resolve the query by resolving the location of the city. The plan recognizes that “san Francisco” is a city or town type of location (e.g., ns:location.citytown) and therefore issues a request for a ?CITYGEO variable defined by a latitude ?CITYLAT and longitude ?CITYLONG. Therefore, a provider can provide the ?CITYGEO variable and its subcomponents in response to the lexical “San Francisco.” Subsequent iterations (not shown) can search for the point of interest (?POI) of City Hall that is near San Francisco, and then from that point of interest, search for disabled/handicapped accessible parking near City Hall.

[00117] Fig. 7D is a diagram 755 illustrating such a subsequent iteration. With the result for the Location, the point of interest (POI) can be resolved. This is performed by a connection being stitched into the request to the POISearch, as further shown by a template illustrated in Fig. 7E.

[00118] Fig. 7E is a diagram 760 illustrating an example embodiment of stitching variables. The stitching template 762 is matched against the query 764. A stitching template is a template of how certain variables are related to each other. Various stitching templates can be available for searching, and in one example can be stored in a database. Each query can be compared to stitching templates to find the best match based on matching variable types and connections between the variables in the query to variable types and connections within each template, and then connections can be copied from the best matching template to the query. Since the pattern of connections and variables in the stitching template 764 matches the fragment of the query 764 starting at the variable ?POI of the query 764, the extra missing connection from the variable ?POI to the variable ?CITYGEO is added.

Therefore, when the POI search of Fig. 7D resolves the search for “city hall,” it bases the location on the ?LONG and ?LAT of ?CITY.

[00119] Fig. 7F is a diagram 765 illustrating a third iteration of constructing the plan for the query 766. With the Location of the city and the location of the POI in hand, the plan can develop a search for the final parking venue. The POI of city hall can be passed to the PID search for parking. This can be performed by stitching the connection to the GEO location of OBJ1 of Fig. 7G to the PID search by matching the template from Fig. 7G to the query 766.

[00120] Fig. 7G is a diagram 770 illustrating an example embodiment of stitching variables. In the stitching template 772, ?OBJ1, having a connection to an ?OBJ2 variable which in turn connects to a ?GEO variable and ?LONG and ?LAT variables, is matched to the ?PID of the parking venue. The connection from ?OBJ1 to ?GEO of the stitching template 772 is stitched to connect ?PID and ?POIGEO in the query 774. Therefore, when the POI of the search of Fig. 7F resolves the search for the parking venue, it bases the location on the ?LONG and ?LAT of ?POI, now connected to the variable ?PID, which in this figure represents the final parking venue.

[00121] Fig. 8 is a diagram 800 illustrating an example embodiment of a method of implementing the intelligent knowledge layer/core employed by the present invention. The method begins (802) and matches input query against provider capability descriptions (804). Then, the method composes matched providers to generate candidate query plans (806). The method then determines whether valid plans are generated, for example, by determining whether the input query is covered by matched provider’s capabilities (808). If not, the method ends (826).

[00122] However, if valid plans are generated (808), the method selects a top query plan and removes plan from the list of candidates (810). The top query plan can be a plan with a lowest cost of using the providers if each provider charges, a plan expected to return a result the fastest based on speed analytics about each provider, etc. If a plan is not selected (812), the method ends (826). However, if a plan is selected (812), the method executes the selected plan, started by invoking the inner most provider of the plan, and applying those results to the outer most provider (814). The inner most provider indicates the provider that does not require results from any other provider. Once the result is provided from the inner most provider, another provider in the chain can provide an answer. If the invoked provider does not return results (816), the method selects the next top query plan and removes that plan

from the candidates (810). However, if it does return results, the method then applies hard constraints and relaxes soft constraints when appropriate (818). Then, if no more providers remain in the plan to execute (820), the method returns the results (824) and ends (826). However, if more providers remain, (820), the method applies the partial results to dependent providers (822) and then continues executing the selected plan by invoking the next inner most provider, using the partial results if necessary (814).

[00123] Fig. 9 illustrates a computer network or similar digital processing environment in which embodiments of the present invention may be implemented.

[00124] Client computer(s)/devices 50 and server computer(s) 60 provide processing, storage, and input/output devices executing application programs and the like. The client computer(s)/devices 50 can also be linked through communications network 70 to other computing devices, including other client devices/processes 50 and server computer(s) 60. The communications network 70 can be part of a remote access network, a global network (e.g., the Internet), a worldwide collection of computers, local area or wide area networks, and gateways that currently use respective protocols (TCP/IP, Bluetooth®, etc.) to communicate with one another. Other electronic device/computer network architectures are suitable.

[00125] Fig. 10 is a diagram of an example internal structure of a computer (e.g., client processor/device 50 or server computers 60) in the computer system of Fig. 9. Each computer 50, 60 contains a system bus 79, where a bus is a set of hardware lines used for data transfer among the components of a computer or processing system. The system bus 79 is essentially a shared conduit that connects different elements of a computer system (e.g., processor, disk storage, memory, input/output ports, network ports, etc.) that enables the transfer of information between the elements. Attached to the system bus 79 is an I/O device interface 82 for connecting various input and output devices (e.g., keyboard, mouse, displays, printers, speakers, etc.) to the computer 50, 60. A network interface 86 allows the computer to connect to various other devices attached to a network (e.g., network 70 of Fig. 9).

Memory 90 provides volatile storage for computer software instructions 92 and data 94 used to implement an embodiment of the present invention (e.g., reasoning interface layer, intelligent knowledge/core layer, handler modules, context mapper, reasoning arbitration, consistency checker, meta reasoner, provider plugins, provider arbitration and planning, query execution and knowledge fusion, constraint satisfaction and relaxation, data inference

code detailed above). Disk storage 95 provides non-volatile storage for computer software instructions 92 and data 94 used to implement an embodiment of the present invention. A central processor unit 84 is also attached to the system bus 79 and provides for the execution of computer instructions.

[00126] In one embodiment, the processor routines 92 and data 94 are a computer program product (generally referenced 92), including a non-transitory computer-readable medium (e.g., a removable storage medium such as one or more DVD-ROM's, CD-ROM's, diskettes, tapes, etc.) that provides at least a portion of the software instructions for the invention system. The computer program product 92 can be installed by any suitable software installation procedure, as is well known in the art. In another embodiment, at least a portion of the software instructions may also be downloaded over a cable communication and/or wireless connection. In other embodiments, the invention programs are a computer program propagated signal product embodied on a propagated signal on a propagation medium (e.g., a radio wave, an infrared wave, a laser wave, a sound wave, or an electrical wave propagated over a global network such as the Internet, or other network(s)). Such carrier medium or signals may be employed to provide at least a portion of the software instructions for the present invention routines/program 92.

[00127] The teachings of all patents, published applications and references cited herein are incorporated by reference in their entirety.

[00128] While example embodiments have been particularly shown and described, it will be understood by those skilled in the art that various changes in form and details may be made therein without departing from the scope of the embodiments encompassed by the appended claims.

What is claimed is:

1. A method of processing queries at an interactive voice response system, the method comprising:
 - determining, based on a received query and contextual information, a plurality of candidate reasoners to respond to at least a portion of the received query ;
 - generating, at each candidate reasoner determined, at least one additional or revised query field or value based on the contextual information, the received query, and a rule of a rule database;
 - merging the at least one additional or revised query field or value for each candidate reasoner based on a confidence score associated with a result of each corresponding candidate reasoner, the confidence score based on applicability of the contextual information to the received query; and
 - providing an enhanced query, the enhanced query having the at least one additional or revised query field or value.
2. The method of Claim 1, further comprising:
 - determining whether the at least one additional or revised query field or value should be input to at least one candidate reasoner to generate further additional or revised query fields or values.
3. The method of Claim 1, further comprising:
 - generating the received query based on a natural language understanding request.
4. The method of Claim 1, wherein a plurality of the candidate reasoners can be configured to be considered, the plurality of candidate reasoners including at least one of a spatial reasoning module configured to perform reasoning based on locations of information in the query, a temporal reasoning module configured to perform reasoning based on timing of information in the query, contextual reasoning configured to perform reasoning based on contextual information , preference reasoning configured to perform reasoning based on preferences of the users, and

custom reasoning configured to perform reasoning based on domain specific information present in the context or received query.

5. The method of Claim 1, further comprising:
 - gathering contextual information from sensors of a system, the system being separate from the interactive voice response system, receiving the query.
6. The method of Claim 1, wherein merging the at least one additional or revised query fields or values further includes prioritizing at least one additional or revised query fields or values from a reasoner ranked with a higher priority.
7. The method of Claim 1, wherein the determining, generating, merging, and providing are performed by a reasoning interface layer, the method further comprising:
 - providing the enhanced query to an intelligent knowledge layer configured to resolve each of the additional query fields of the enhanced query.
8. The method of Claim 1, further comprising:
 - determining, based on the enhanced query a plurality of sources, a source being an external content provider, a virtual assistant, or internal database, that can be used to fulfill all or a portion of the enhanced query.
 - constructing a plan, the plan indicating which sources to invoke and the order to invoke the sources.
9. A system for processing queries at an interactive voice response system, the system comprising:
 - a processor; and
 - a memory with computer code instructions stored therein, the memory operatively coupled to said processor such that the computer code instructions configure the processor to implement a reasoning interface layer that is configured to:
 - determine, based on a received query and contextual information, a plurality of candidate reasoners to respond to at least a portion of the received query;

generate, at each candidate reasoner determined, at least one additional or revised query field or value based on the contextual information and a rule of a rule database;

merge the at least one additional or revised query field or value for each candidate reasoner based on a confidence score associated with a result of each corresponding candidate reasoner, the confidence score based on applicability of the contextual information to the received query; and

provide an enhanced query, the enhanced query having the at least one additional or revised query field or value.

10. The system of Claim 9, wherein the reasoning interface layer is further configured to:
determine whether the at least one additional or revised query field or value should be input to at least one candidate reasoner to generate further additional or revised query fields or values.
11. The system of Claim 9, wherein the reasoning interface layer is further configured to:
generate the received query based on a natural language understanding request.
12. The system of Claim 9, wherein a plurality of the candidate reasoners can be configured to be considered, the plurality of candidate reasoners including at least one of a spatial reasoning module configured to perform reasoning based on locations of information in the query, a temporal reasoning module configured to perform reasoning based on timing of information in the query, contextual reasoning configured to perform reasoning based on contextual information, preference reasoning configured to perform reasoning based on preferences of the users, and custom reasoning configured to perform reasoning based on domain specific information present in the context or received query.
13. The system of Claim 9, wherein the reasoning interface layer is further configured to:
receive contextual information from sensors of a system, the system separate from the interactive voice response system, receiving the query.

14. The system of Claim 9, wherein merging the at least one additional or revised query fields or values further includes prioritizing at least one additional or revised query fields or values from a reasoner ranked with a higher priority.
15. The system of Claim 9, wherein the reasoning interface layer is further configured to:
 - providing the enhanced query to an intelligent knowledge layer configured to resolve each of the additional query fields of the enhanced query.
16. The system of Claim 9, wherein the reasoning interface layer is further configured to:
 - determine, based on the enhanced query a plurality of sources, a source being an external content provider, a virtual assistant, or internal database, that can be used to fulfill all or a portion of the enhanced query.
 - construct a plan, the plan indicating which sources to invoke and the order to invoke the sources.
17. A non-transitory computer-readable medium configured to store instructions for processing queries at an interactive voice response system, the instructions, when loaded and executed by a processor, causes the processor to:
 - determine, based on a received query and contextual information, a plurality of candidate reasoners to respond to at least a portion of the received query;
 - generate, at each candidate reasoner determined, at least one additional or revised query field or value based on the contextual information, the received query, at each candidate reasoner to a rule of a rule database;
 - merge the at least one additional or revised query field or value for each candidate reasoner based on a confidence score associated with a result of each corresponding candidate reasoner, the confidence score based on applicability of the contextual information to the received query; and
 - provide an enhanced query, the enhanced query having the at least one additional or revised query field or value.
18. The non-transitory computer readable medium of Claim 15, wherein the instructions further cause the processor to:

determine whether the at least one additional or revised query field or value should be input to at least one candidate reasoner to generate further additional or revised query fields or values.

19. The non-transitory computer readable medium of Claim 15, wherein the instructions further cause the processor to:
 - generate the received query based on a natural language understanding request.
20. The non-transitory computer readable medium of Claim 15, wherein a plurality of the candidate reasoners can be configured to be considered, the plurality of candidate reasoners including include at least one of a spatial reasoning module configured to perform reasoning based on locations of information in the query, a temporal reasoning module configured to perform reasoning based on timing of information in the query, contextual reasoning configured to perform reasoning based on contextual information, preference reasoning configured to perform reasoning based on preferences of the users, and custom reasoning configured to perform reasoning based on domain specific information present in the context or received query.
21. A method comprising:
 - matching portions of a received query to corresponding values in a database of provider capabilities;
 - determining, for each portion of the received query, a respective provider to execute the portion of the received query based on the provider capabilities;
 - sending each portion of the received query to its respective provider; and
 - combining results from each respective provider to form a returned query result.
22. The method of Claim 21, wherein combining the results further includes, based on the results from each provider, filtering a list of results by applying hard constraints from portions of the received query, and ranking a list of results based on a priority of soft constraints from portions of the received query.

23. The method of Claim 21, further comprising:
 - determining whether a first respective provider requires receiving results from a second provider, and, if so, sending appropriate results of the second provider to the first provider as input.
24. The method of Claim 21, wherein the respective providers include at least one of the following: a third party content provider application programming interface (API), an internal database, and a virtual assistant.
25. The method of Claim 21, further comprising providing a semantic description corresponding to each respective provider, wherein determining a respective provider to execute the portion of the received query includes determining whether the portion of the received query includes a request for a resource described by the semantic description corresponding to the respective provider.
26. The method of Claim 21, further comprising constructing a plan that includes an order of invoking each respective provider.
27. The method of Claim 21, wherein combining results from each respective provider to return a query result further includes:
 - determining a first result from a first provider; and
 - sending the first result to a second provider as part of the portion of the query to the second provider;
 - wherein the first result from the first provider is employed to resolve the portion of the query to the second provider.
28. The method of Claim 21, further comprising relaxing at least one portion of the query if a combination of fields or values of the query produces an incompatible result.
29. The method of Claim 21, further comprising:
 - comparing variables and connections between variables of a received query to variables and connections between variables to at least one stitching template to find a best matching stitching template;

finding a connection present in the best matching stitching template from a first template variable to a second template variable that is absent in the portion of the received query between a first query variable and a second query variable, the first query variable and second query variable corresponding to the first template variable and second template variable; and

copying at least one connection from the stitching template to the portion of the query between the first query variable and second query variable.

30. A system comprising:

a processor; and

a memory with computer code instructions stored therein, the memory operatively coupled to said processor such that the computer code instructions configure the processor to implement:

a planning module configured to match a portion of a received query to corresponding values in a database of provider capabilities, and determine, for each portion of the received query, a respective provider to execute the portion of the received query based on the provider capabilities;

a query execution module configured to send each portion of the received query to its respective provider, and combine results from each respective provider to form a returned query result.

31. The system of Claim 30, wherein combining the results further includes, based on the results from each provider, filtering a list of results by applying hard constraints from portions of the received query, and ranking a list of results based on a priority of soft constraints from portions of the received query.

32. The system of Claim 30, wherein:


the planning module is further configured to determine whether a first provider requires receiving results from a second provider, and,

the query execution module is configured to, if the first provider requires receiving results from the second provider, send appropriate results of the second provider to the first provider as input.

33. The system of Claim 30, wherein the respective providers include at least one of the following: a third party content provider application programming interface (API), an internal database, and a virtual assistant.
34. The system of Claim 30, wherein the planning module is provided a semantic description corresponding to each respective provider, and determining a respective provider to execute the portion of the received query includes determining whether the portion of the received query includes a request for a resource described by the semantic description corresponding to the respective provider.
35. The system of Claim 30, wherein the planning module is further configured to construct a plan that includes an order of invoking each respective provider.
36. The system of Claim 30, wherein the query execution module is further configured to combine results from each respective provider to return a query result by:
 - determining a first result from a first provider; and
 - sending the first result to a second provider as part of the portion of the query to the second provider;
 - wherein the first result from the first provider is employed to resolve the portion of the query to the second provider.
37. The system of Claim 30, further comprising constraint relaxation module configured to relax at least one portion of the query if a combination of fields or values of the query produces an incompatible result.
38. The system of Claim 30, wherein the planning module is further configured to:
 - compare variables and connections between variables of a received query to variables and connections between variables to at least one stitching template to find a best matching stitching template;
 - find a connection present in the best matching stitching template from a first template variable to a second template variable that is absent in the portion of the received query between a first query variable and a second query variable, the first query variable and second query variable corresponding to the first template variable and second template variable; and

copy at least one connection from the stitching template to the portion of the query between the first query variable and second query variable.

Fig. 1 100







Fig. 2 200

Fig. 3 300



400




Fig. 4

Fig. 5 500




Fig. 6

600





Fig. 7B

STEP 2: Identify Candidate Plugins



?PID → Parking
?POI → POISearch
?CITY → Location

Fig. 7A

STEP 1: Send Query to IKA

"Find disabled parking near city hall in san francisco."

```

select ?PID ?NAME ?PRICE where {
  ?PID ns:type.object.type nuan:parking.generic .
  ?PID ns:type.object.name ?NAME .
  ?PID nuan:base.entity.amentities are ?AMENTITIES .
  ?PID nuan:parking.generic.intended_for occupant ?OCCUPANT .
  ?PID nuan:cost.quantitative.typical_cost ?PRICE .
  ?POI ns:type.object.type ns:travel.tourist_attraction .
  ?CITY ns:type.object.type ns:location.citytown .
  ?PID ns:location.location.is_near ?POI .
  ?POI ns:type.object.lexical "city hall" .
  ?POI ns:location.location.street_address ?ADDR .
  ?ADDR ns:location.mailing_address.citytown ?CITY .
  ?CITY ns:type.object.lexical "san francisco" .
  FILTER SOFT (?AMENTITIES = nuan:amenity.elevator)
  FILTER (?OCCUPANT = ns:person.disabled_person)
  ORDER BY ASC(?PRICE)
}

```

Query 710

Construct Plan




Fig. 7C 750





Fig. 7E 760

Stitches

Stitching Template 772



Query 774





Fig. 7G 770

Fig. 8

800






FIG. 9




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/055860

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06N5/04
ADD. G06F17/27 G06F16/33

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06N G06F

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	TAKAHISA FUJINO ET AL: "A SPARQL Query Rewriting Approach on Heterogeneous Ontologies with Mapping Reliability", 2012 IIAI INTERNATIONAL CONFERENCE ON ADVANCED APPLIED INFORMATICS, 1 September 2012 (2012-09-01), pages 230-235, XP055546129, DOI: 10.1109/IIAI-AAI.2012.54 ISBN: 978-0-7695-4826-5 page 230 - page 235, left-hand column, last paragraph	1-38
A	WO 2006/014824 A1 (WIRELESS 5TH DIMENSIONAL NETWO [US]; FULLER NAPIER S [US]) 9 February 2006 (2006-02-09) the whole document ----- -/-	1-38



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

23 January 2019

Date of mailing of the international search report

04/02/2019

Name and mailing address of the ISA/

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

Authorized officer

Volkmer, Markus

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/055860

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2007/124263 A1 (KATARIYA SANJEEV [US] ET AL) 31 May 2007 (2007-05-31) the whole document -----	1-38
A	US 2011/087629 A1 (B FAR REZA [US] ET AL) 14 April 2011 (2011-04-14) the whole document -----	1-38
A	EP 2 648 137 A2 (ERICSSON TELEFON AB L M [SE]) 9 October 2013 (2013-10-09) the whole document -----	1-38
A	Carlos R. Rivero ET AL: "A Reference Architecture for Building Semantic-Web Mediators" In: "Serious Games", 1 January 2011 (2011-01-01), Springer International Publishing, Cham 032682, XP055546081, ISSN: 0302-9743 ISBN: 978-3-642-37803-4 vol. 8827, pages 330-341, DOI: 10.1007/978-3-642-22056-2_36, the whole document -----	1-38
X,P	Prateek Jain ET AL: "Nuance Reasoning Framework: A Rule-Based System for Semantic Query Rewriting : Second International Joint Conference, RuleML+RR 2018, Luxembourg, Luxembourg, September 18-21, 2018, Proceedings" In: "Serious Games", 1 January 2018 (2018-01-01), Springer International Publishing, Cham 032682, XP055544857, ISSN: 0302-9743 ISBN: 978-3-642-37803-4 vol. 11092, pages 285-292, DOI: 10.1007/978-3-319-99906-7_20, the whole document -----	1-38

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2018/055860

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2006014824 A1	09-02-2006	EP 1779269 A1 JP 2008507792 A US 2007198505 A1 WO 2006014824 A1	02-05-2007 13-03-2008 23-08-2007 09-02-2006
US 2007124263 A1	31-05-2007	CN 101356525 A KR 20080072673 A US 2007124263 A1 WO 2007064482 A1	28-01-2009 06-08-2008 31-05-2007 07-06-2007
US 2011087629 A1	14-04-2011	US 2011087629 A1 US 2011093469 A1	14-04-2011 21-04-2011
EP 2648137 A2	09-10-2013	EP 2648137 A2 US 2013262366 A1	09-10-2013 03-10-2013