

SYSTEM AND METHOD FOR ANALYSIS OF DATA FROM MULTIPLE DATA SOURCES

Related Cases

[001] This application claims the benefit of U.S. Provisional Application No. 63/134,677, filed on 07 January 2021, the contents of which is incorporated herein by reference.

Background

[002] Data analytics is impacting businesses worldwide. Applying data analytics can lead to discovery of hidden correlations in data which can then be used by businesses. However, most analytics solutions available in the market tend to focus on how to load a customer's data into analytics tools and analyze the data from the analytics interface. Therefore, the target users for these solutions may be expected to have some data science background. On the other hand, regular users are more familiar with legacy data query interfaces (e.g., a reporting center with a fixed list of reports). However, a legacy data query interfaces will usually highlight ad-hoc query instead of context-based conversational query which could benefit from analytic results.

Brief Summary of Disclosure

[003] The disclosure may include but is not limited to a computing system for analysis of data from multiple data sources, a computer implemented method for analysis of data from multiple data sources, a computer program product residing on a computer readable storage medium, and a computing system including one or more processors and one or more memories.

[004] In an example implementation, a computer-implemented method for analysis of data from multiple data sources may include receiving a query instance. The query instance may be processed to generate one or more relational query instances based, at least in part, on the query instance. The generated one or more relational query instances may

be distributed among one or more nodes. Each of the generated one or more relational query instances may be processed by at least one node of the one or more nodes to extract data from at least one data source of multiple data sources in communication therewith corresponding to the one or more relational query instances. One or more insights may be generated based on the extracted data. The one or more insights generated may be stored in relevant cache to be retrieved therefrom.

[005] One or more of the following example features may be included. Entities in the query instance and attributes associated therewith may be identified. Historic data from at least one data source of the multiple data sources related to one or more of the entities and attributes identified may be retrieved. At least one model may be trained based on the historic data retrieved. Tags may be associated with the at least one trained model indicative of corresponding one or more of the entities and attributes. The at least one trained model with associated tags may be stored in at least one data source of the multiple data sources, to be retrieved therefrom. The at least one data source of the multiple data sources may be monitored for updates in the historic data. Entities and attributes associated with the updates in the historic data may be determined. The at least one model, with associated tags related to the determined entities and attributes, may be re-trained based on the updates in the historic data. Monitoring the at least one data source of the multiple data sources may further include periodically monitoring the at least one data source of the multiple data sources for updates in the historic data. The generated one or more insights may be associated to the corresponding query instance received. The query instance received, with the one or more insights associated therewith, may be cached in an enterprise query base (EQB) system, to be retrieved therefrom. A query context may be determined for the query instance. The cached query instance may be linked to the query context based on a matching algorithm. The matching algorithm may be a semantic matching algorithm. Entities in the query instance and attributes associated therewith may be identified. The query context may be determined from the identified entities and attributes therefrom. One

or more predefined parameters may be filtered from the query instance. A hash value for the query instance received may be generated. The hash value may be used to find a query cache that matches the hash value. A query plan based on the processing of the query instance may be generated. Another related query instance may be processed based on the query plan. The query instance may be a query text with literal parameters filtered. The extracted data may be processed to generate a response based on the query instance.

[006] In another example implementation, a computing system for analysis of data from multiple data sources may include an enterprise query base configured to communicate with at least one data source of the multiple data sources to access data therefrom. The enterprise query base may be configured to receive a query instance. The enterprise query base may be further configured to process the query instance to generate one or more relational query instances based upon, at least in part, the query instance, with each of the one or more relational query instances to be processed by at least one node of one or more nodes. The one or more relational query instances may be distributed among the one or more nodes to extract data from the at least one data source in communication therewith corresponding to the one or more relational query instances. The enterprise query base may be further configured to receive extracted data from each of the one or more nodes queried. The computing system may further include an analytic service in communication with the enterprise query base. The analytic service may be configured to generate one or more insights based on the extracted data. The analytic service may be further configured to store the one or more insights generated in at least one data source of the multiple data sources, to be retrieved therefrom.

[007] One or more of the following example features may be included. The enterprise query base may be further configured to identify entities in the query instance and attributes associated therewith. The enterprise query base may be further configured to: retrieve historic data from the at least one data source of the multiple data sources related to one or more of the entities and attributes identified; train at least one model based on the historic

data retrieved; associate tags with the at least one trained model indicative of corresponding one or more of the entities and attributes; and store the at least one trained model with associated tags in at least one data source of the multiple data sources, to be retrieved therefrom. The enterprise query base may be further configured to: monitor the at least one data source of the multiple data sources for updates in the historic data; determine entities and attributes associated with the updates in the historic data; and re-train the at least one model, with associated tags related to the determined entities and attributes, based on the updates in the historic data. The enterprise query base monitoring the at least one data source of the multiple data sources may further include the enterprise query base configured to periodically monitor the at least one data source of the multiple data sources for updates in the historic data. The enterprise query base may be further configured to: associate the generated one or more insights to the corresponding query instance received; and cache the query instance received, with the one or more insights associated therewith, in at least one data source of the multiple data sources, to be retrieved therefrom. The enterprise query base may be further configured to determine a query context for the query instance and may link the cached query instance to the determined query context based on a matching algorithm. The matching algorithm may be a semantic matching algorithm. The query context may be determined from identified entities in the query instance and attributes associated therewith. The enterprise query base may be further configured to filter one or more predefined parameters from the query instance. The enterprise query base may be further configured to generate a hash value for the query instance received. The enterprise query base may be configured to use the hash value to find a query cache that most closely matches the hash value. The enterprise query base may be further configured to generate a query plan based on the processing of the query instance. The enterprise query base may be further configured to process another related query instance based on the query plan. The query instance may be a query text with literal parameters filtered. The one or more nodes may be arranged in a hierarchical structure. The one or more relational query

instances may be distributed among the one or more nodes directionally from upper nodes to lower nodes. The enterprise query base may be further configured to process the extracted data to generate a response based on the query instance.

[008] In an example implementation, a computer-implemented method for analysis of data from multiple data sources may include receiving a query instance. The query instance may be processed to generate one or more relational query instances based upon, at least in part, the query instance. Each of the generated one or more relational query instances may be processed by at least one node of one or more nodes. The one or more relational query instances may be distributed among the one or more nodes to extract data from the at least one data source. The extracted data may be received from each of the one or more nodes queried. One or more insights may be generated based on the extracted data. The one or more insights generated may be stored in at least one data source of the multiple data sources, to be retrieved therefrom.

[009] One or more of the following example features may be included. Entities in the query instance and attributes associated therewith may be identified. Historic data from at least one data source of the multiple data sources related to one or more of the entities and attributes identified may be retrieved. At least one model may be trained based on the historic data retrieved. Tags may be associated with the at least one trained model indicative of corresponding one or more of the entities and attributes. The at least one trained model with associated tags may be stored in at least one data source of the multiple data sources, to be retrieved therefrom. The at least one data source of the multiple data sources may be monitored for updates in the historic data. Entities and attributes associated with the updates in the historic data may be determined. The at least one model, with associated tags related to the determined entities and attributes, may be re-trained based on the updates in the historic data. The generated one or more insights may be associated to the corresponding query instance received. The query instance received, with the one or more insights associated therewith, may be cached in an enterprise query base (EQB) system, to

be retrieved therefrom.

[0010] In yet another example implementation, a computer program product residing on a computer readable storage medium having a plurality of instructions stored thereon is provided. The computer program product, when executed across one or more processors, causes at least a portion of the one or more processors may perform operations comprising receiving a query instance. The query instance may be processed to generate one or more relational query instances based at least in part on the query instance. Each of the generated one or more relational query instances being processed by at least one node of one or more nodes. The one or more relational query instances may be distributed among the one or more nodes to extract data from the at least one data source. The extracted data may be received from each of the one or more nodes queried. One or more insights may be generated based on the extracted data. The one or more insights generated may be stored in at least one data source of multiple data sources, to be retrieved therefrom.

[0011] One or more of the following example features may be included. Entities in the query instance and attributes associated therewith may be identified.

[0012] Other example implementations or embodiments may include: a system as shown and described in this disclosure, a query base system (e.g., enterprise query base), an enterprise query base node, a hierarchy of enterprise query base nodes, a method for analysis of data from multiple data sources as shown and described in this disclosure, a method for implementing an enterprise query base system, a computer-implemented method substantially as described with reference to any of the examples and/or to any of the drawings in this disclosure, a computing system including one or more processors and one or more memories configured to perform operations substantially as described with reference to any of the examples and/or to any of the drawings in this disclosure, a computer program product residing on a computer readable storage medium having a plurality of instructions stored thereon which, when executed across one or more processors, causes at least a portion of the one or more processors to perform operations

substantially as described with reference to any of the examples and/or to any of the drawings of this disclosure, and an apparatus configured substantially as described in this disclosure with reference to any of the examples and/or to any of the drawings in this disclosure.

[0013] The details of one or more example implementations are set forth in the accompanying drawings and the description below. Other possible example embodiments, features, aspects, and/or possible example advantages will become apparent from the description, the drawings, and the claims. Some implementations may not have those possible example features and/or possible example advantages, and such possible example embodiments, features, aspects, and/or possible example advantages may not necessarily be required of some implementations.

Brief Description of the Drawings

[0014] FIG. 1 is an example diagrammatic view of a query base process coupled to an example distributed computing network according to one or more example implementations of the disclosure;

[0015] FIG. 2 is an example diagrammatic view of a client electronic device of FIG. 1 according to one or more example implementations of the disclosure;

[0016] FIG. 3 is an example diagrammatic view of a network environment implementing a query base system with multiple enterprise query base nodes distributed therein according to one or more example implementations of the disclosure;

[0017] FIG. 4 is an example diagrammatic view of a query base system including a query base subsystem in communication with multiple data sources therein according to one or more example implementations of the disclosure;

[0018] FIG. 5A is an example detailed diagrammatic view of a query base subsystem having a master client node and a server node according to one or more example implementations of the disclosure;

[0019] FIG. 5B is an example detailed diagrammatic view of an enterprise query base

node having a master client module and a server module according to one or more example implementations of the disclosure;

[0020] FIG. 6 is an example diagrammatic view of a proxy table of a server node being mapped to a base table associated with a remote data source according to one or more example implementations of the disclosure;

[0021] FIG. 7 is another example diagrammatic view of the enterprise query base node propagating a query therein according to one or more example implementations of the disclosure;

[0022] FIG. 8 is an example diagrammatic view of a system for analysis of data from multiple data sources according to one or more example implementations of the disclosure;

[0023] FIG. 9 is an example diagrammatic view of a query cache according to one or more example implementations of the disclosure;

[0024] FIG. 10 is an example flow diagram of a query in a query cache according to one or more example implementations of the disclosure;

[0025] FIG. 11 is an example schematic of an entity model according to one or more example implementations of the disclosure;

[0026] FIG. 12 is an example flowchart of the query base process for analysis of data from multiple data sources according to one or more example implementations of the disclosure; and

[0027] FIG. 13 is another example flowchart of the query base process for analysis of data from multiple data sources according to one or more example implementations of the disclosure.

[0028] Like reference symbols in the various drawings indicate like elements.

Detailed Description

[0029] As noted above, there are many analytics tools offered in the market, such as Microsoft™ Power BI, SAP™ Analytics, Looker™, and Tableau™ Server, etc. The key

features of these analytics tools may include loading data from different sources, providing interfaces for data training, and visualization of descriptive data as well as analytic results, etc. However, the user of these tools may need to understand data modeling and analytic concepts in order to build a visualized interface. The result is usually published before it may be consumed by normal audiences. For example, Microsoft™ Power BI may be used to create a report. The Microsoft™ Power BI may be loaded with data (e.g., from an Excel™ file or from a data system). A pop-up order may be chosen regarding which kind of visualization tools may be needed to use e.g., chart or pie chart based on data. Then, the Microsoft™ system may use a Power BI service to publish the report and then subscribe to what may be published (e.g., email may be sent indicating that the report is available).

[0030] These tools may need offline data training and modelling to be deployed for use. This may be typically slow and ineffective. Also, transactional data may not be linked to insights found by typical data analysis. The target users for these analytics tools may be expected to have some data science background. Further, legacy data query system may usually highlight ad-hoc query instead of context-based conversational query which may benefit from analytic results. Overall, there may be a general quality problem for some analytics tools with respect to insights and suggestions where these insights/suggestions may or may not be relevant depending on the user and/or company.

[0031] Moreover, businesses face a growing problem of data proliferation. Data proliferation may involve information typically being generated and stored in multiple data sources across, for example, an enterprise. Systems, such as database systems, managing large amounts of data may distribute and/or replicate that data across multiple data sources, often in different locations, including on premise, remotely and/or on the cloud, for any of a number of reasons, including, e.g., security issues, disaster prevention and recovery issues, data locality and availability issues, etc. Such database systems may be implemented either as a distributed database or a distributed file system that may tend not to scale well for data mining and business intelligence applications that may require fast

and efficient retrieval and processing of large volumes of data. Furthermore, such database systems may pose problems for third party applications, which may face the challenge of accessing data behind a firewall without copying or replicating the sensitive data from the enterprise data source. One example technique may be to copy or replicate data from one system into another system, when there may be a need to utilize information that is present in the other system. This is typically done using Extract, Transform, and Load (ETL) technology or replication technology. Some enterprises may implement enterprise data hubs (EDHs), which may provide a more practical and scalable solution than other enterprise data systems. An enterprise data hub may be a large storage repository that holds a vast amount of raw data in its native format until it is needed for enterprise-wide information storage and sharing. The enterprise data hubs may provide that all content from disparate sources ingested into the enterprise data hub and made available for search. While some applications may be best served by copying all the data into a data warehouse or a data lake, it may be beneficial to minimize data redundancy and increase the freshness of data by providing ways to access the data from its native location without the need of moving the data to a single data warehouse or a single data lake. Moreover, in some cases, legal aspects may forbid copying sensitive data to an external data warehouse or data lake. Some known solutions may use a common technique to build a data cube or data warehouse for analytics in advance.

[0032] The disclosure may relate to a computing system to address the example and non-limiting data proliferation and analytics issue, for example, in an enterprise environment, where information may be typically generated and stored in multiple data sources. The disclosure may provide an enterprise query base process and/or other processes. The enterprise query base process and/or other processes of the disclosure may integrate the data analytics as part of query engines which may process data from different data sources (e.g., link insights with transactions data). The insights may help to explain the query result from live data, and may link relevant factors to justify an answer. The

disclosed system and method or process may embed common data modeling with an Enterprise Query Broker (EQB) (also referred to as an “enterprise query database”) which may not only connect to different data sources but may also join the data across different data sources. The user may not need to understand the physical data storage issue as a database expert and may use a data entity (or data entities) in a semantic manner. Herein, the insights, analytic models, and results may be used by a user without waiting for publication since those may be linked to related user queries regularly and automatically based on relevant semantics and correlations. In an example implementation, a user may ask a question about a topic (e.g., how many customers go to a restaurant). The disclosed system may use a forecast model to answer this question as well as use insights (e.g., based on insight history) to provide suggestions and correlations. By linking these correlations/suggestions with the answer, the user may find that reported analytics may be particularly relevant to them and decide to act on these analytics (e.g., add promotions during specific holidays especially increase customers). Moreover, data federation and open data port may provide a way to access integrated data from multiple sources including analytic results as well as a data source. Analytic results may provide additional dimensions for an intelligent query answering engine. The disclosed system and method may provide a mechanism to integrate live data with analytic data and analytic models in the enterprise query system.

[0033] The disclosed EQB may cover data analytics results as a kind of data source. The EQB may not only visualize descriptive data but may also answer questions from data analytics such as determining correlations and insights from the data as well as providing forecast/prediction with analytic models based on context. The insights, correlated probabilities, and time series forecast models may then be saved into an insight repository and trained models database. The insight repository and models database may be connected as a data source for the EQB. The EQB may be able to integrate the insights with other live data query answers, and may also make some on-the-fly calculations with probability

insights as well as generate conditional probability for the user query. Forecast models may be joined with the transactional data and answer real time prediction related user queries. This may significantly reduce the released cycle of newly found analytic insights and update analytic forecast models dynamically. The EQB may support query text cache which may link correlated queries and execution plans of the query. This may provide an infrastructure to run multiple correlated queries in a dialogue-based user query system. Further, the disclosed analytic service (e.g., also referred to as analytic tool, analytics tool, analytics service) may be configured to pull data through the EQB for data explore and analysis. The analytic service may enhance the ad-hoc prediction capability of the query system and provide more accurate analytics thus more accurate suggestions.

[0034] In some example implementations, the disclosed system may make use of a query communication mesh concept to provide a way to access and process integrated data from multiple data sources in a network without creating redundant data stores. This may provide a way to distribute and push down the query to the data stored at its native location and avoid the need of transferring large amounts of sensitive data across the network. The enterprise query base process and/or other processes of the disclosure may implement an elastic mesh technique to tackle various complex enterprise data source issues which may include hybrid data sources including on premise and in cloud, and/or systems across geographic regions, different firewalls, and different networks separated by firewalls, etc.

[0035] In some example implementations, the system may be implemented as a directional graph or tree-like mesh that passes the query from an application client to a data source by employing one or more enterprise query base (EQB) nodes. Each of the EQB nodes may include or may either act as one of a client node that receives the query and a server node which, in turn, is connected to one or more back-end data sources to access data therefrom. The client node of the EQB subsystem at the root of the tree may receive the query (often in natural language), and may translate and pass a processed query to the server node (e.g., location of data source). The processed query may be divided into

multiple sub-queries, according to a build plan or an execution plan, to be processed by one or more data sources in communication with the server node of the same EQB subsystem and/or to be selectively distributed among server nodes of some of the other EQB nodes.

[0036] In some example implementations, the query may be pushed down to the data source(s). There may be a clear direction and client/server relationship such as a clear direction from the client node which may receive a natural language query and may process/translate natural language to be distributed among different nodes from processing thereof. The system may send out one query that may move down chain towards different EQB nodes that are able to pass the query in a direction or directions towards multiple data sources that may be associated therewith and have information needed for processing the query. The processed query may not include the calculated route information but instead, each EQB node may keep destination data source mapping and may optimize the query plan itself. That is, each EQB node itself may be a sub-query engine and may process the sub-query (e.g., filtering, joining, aggregating, etc.).

[0037] Distributed databases for large volumes of data, perhaps on the order of terabytes, may be implemented across several servers, each designed to host a portion of a database and typically storing a particular table data. Typically, such distributed database systems may involve a partitioned database or database grid under the same network (e.g., intranet). The objective of the distribution of the data partition in different data sources in existing systems may be typically for performance e.g., using parallel or distributed approaches to access different data partitions concurrently and more efficiently. In some examples, a network of connected nodes may be provided in the form of a distributed federated database system, which is a type of metadata base management system (DBMS), which may transparently map multiple autonomous database systems into a single federated database. The constituent databases may be distributed and interconnected via a computer network and may be geographically decentralized. In traditional database

systems, all nodes may be configured to be able to send/receive queries with all nodes being treated equal (e.g., no client/server relationship). These types of database systems may be generally not able to handle different types of data sources from multiple vendors, e.g., where some data sources are behind a firewall, some data sources are in the cloud, etc. On the other hand, third party applications may also face the challenge of accessing data behind a firewall without copying or replicating the sensitive data from the enterprise data source.

[0038] The disclosed system (also, sometimes, referred to as “query base system”) which may be based on a mesh concept, that may work as a gateway for a data source federation, may distribute queries, and may aggregate the result from different data sources. The query base system may not be a database management system and may not store data with itself, but may be a query engine, which may be used to query different types of databases. Therefore, the system may generally be placed at a level above in the network hierarchy in relation to database management systems. In some example implementations, the system may facilitate connection to different data sources, including different locations of the data source and/or different types of data sources. The system may provide the ability to understand different types of data sources (e.g., Oracle™, MongoDB, IBM™, SAP™, or other data sources). Whereas other data systems may describe querying data within the same organization, the disclosed system may be enabled to query at a level above the organization such as querying different network database systems. Each of these networks may have varying security concerns depending on the network. When the query base system queries the network, the actual network may perform sub-queries of their actual data and then the network may output results to the query base system, thus mitigating security and third-party access concerns. The results obtained from one or more networks may provide the relevant data by merging results together, in contradistinction to other example systems that obtain results from the actual data within their own network.

[0039] In some example implementations, the query base system may use different

operators (e.g., link multiple operators together) for querying data from different data sources instead of simply storing data. The system may cross different networks including the Internet to access one or more data sources behind a firewall. Further, the system may integrate with natural language processing (NLP), which may handle semantic mapping and may remove the barrier for the application to query different data sources with different query languages. The query base system may use semantic operations that may be adapted for different types of data sources, for example, depending on type of data source, the network/geographic location of the data source, available access type, etc. (e.g., allowing the system to access different types of data sources that may require different accessibility). This means that the system may be able to generate a sub-query for one data source that may be completely different from a sub-query generated for a second different data source, thus enabling the system to function across different types of networks and also different types of data source.

System Overview:

[0040] In some example implementations, the disclosure may be embodied as a system, method, apparatus, or computer program product. Accordingly, in some example implementations, the system, method, apparatus, or computer program product of the disclosure may take the form of an entirely hardware implementation, an entirely software implementation (including firmware, resident software, micro-code, etc.) or an implementation combining software and hardware aspects that may all generally be referred to herein as a “circuit,” “module” or “system”. Furthermore, in some example implementations, the disclosure may include a computer program product on a computer-usable storage medium having computer-usable program code embodied in the medium.

[0041] In some example implementations, any suitable computer usable or computer readable medium (or media) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. The computer-usable, or computer-readable, storage medium (including a storage device associated with a computing device or client electronic device) may be, for example, but is not limited to,

an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer-readable medium may include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a digital versatile disk (DVD), a static random access memory (SRAM), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, a media such as those supporting the internet or an intranet, or a magnetic storage device. Note that the computer-usable or computer-readable medium could even be a suitable medium upon which the program is stored, scanned, compiled, interpreted, or otherwise processed in a suitable manner, if necessary, and then stored in a computer memory. In the context of the disclosure, a computer-usable or computer-readable, storage medium may be any tangible medium that can contain or store a program for use by or in connection with the instruction execution system, apparatus, or device.

[0042] In some example implementations, a computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. In some example implementations, such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. In some example implementations, the computer readable program code may be transmitted using any appropriate medium, including but not limited to the internet, wireline, optical fiber cable, RF, etc. In some example implementations, a computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

[0043] In some example implementations, computer program code for carrying out operations of the disclosure may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Java®, Smalltalk, C++ or the like. Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates. However, the computer program code for carrying out operations of the disclosure may also be written in conventional procedural programming languages, such as the "C" programming language, PASCAL, or similar programming languages, as well as in scripting languages such as JavaScript, PERL, or Python. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the internet using an Internet Service Provider). In some example implementations, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGAs) or other hardware accelerators, micro-controller units (MCUs), or programmable logic arrays (PLAs) may execute the computer readable program instructions/code by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the disclosure.

[0044] In some example implementations, the flowchart and block diagrams in the figures show the architecture, functionality, and operation of possible implementations of apparatus (systems), methods and computer program products according to various implementations of the disclosure. Each block in the flowchart and/or block diagrams, and combinations of blocks in the flowchart and/or block diagrams, may represent a module,

segment, or portion of code, which comprises one or more executable computer program instructions for implementing the specified logical function(s)/act(s). These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the computer program instructions, which may execute via the processor of the computer or other programmable data processing apparatus, create the ability to implement one or more of the functions/acts specified in the flowchart and/or block diagram block or blocks or combinations thereof. It should be noted that, in some example implementations, the functions noted in the block(s) may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved.

[0045] In some example implementations, these computer program instructions may also be stored in a computer-readable memory that can direct a computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable memory produce an article of manufacture including instruction means which implement the function/act specified in the flowchart and/or block diagram block or blocks or combinations thereof.

[0046] In some example implementations, the computer program instructions may also be loaded onto a computer or other programmable data processing apparatus to cause a series of operational steps to be performed (not necessarily in a particular order) on the computer or other programmable apparatus to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide steps for implementing the functions/acts (not necessarily in a particular order) specified in the flowchart and/or block diagram block or blocks or combinations thereof.

[0047] Referring now to the example implementation of FIG. 1, there is shown a query base process 10 that may reside on and may be executed by a computer (e.g., computer

12), which may be connected to a network (e.g., network 14) (e.g., the internet or a local area network). Examples of computer 12 (and/or one or more of the client electronic devices noted below) may include, but are not limited to, a personal computer(s), a laptop computer(s), mobile computing device(s), a server computer, a series of server computers, a mainframe computer(s), or a computing cloud(s). In some example implementations, each of the aforementioned may be generally described as a computing device. In certain example implementations, a computing device may be a physical or virtual device. In many example implementations, a computing device may be any device capable of performing operations, such as a dedicated processor, a portion of a processor, a virtual processor, a portion of a virtual processor, portion of a virtual device, or a virtual device. In some example implementations, a processor may be a physical processor or a virtual processor. In some example implementations, a virtual processor may correspond to one or more parts of one or more physical processors. In some example implementations, the instructions/logic may be distributed and executed across one or more processors, virtual or physical, to execute the instructions/logic. Computer 12 may execute an operating system, for example, but not limited to, Microsoft® Windows®; Mac® OS X®; Red Hat® Linux®, or a custom operating system. (Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States, other countries or both; Mac and OS X are registered trademarks of Apple Inc. in the United States, other countries or both; Red Hat is a registered trademark of Red Hat Corporation in the United States, other countries or both; and Linux is a registered trademark of Linus Torvalds in the United States, other countries or both).

[0048] In some example implementations, as will be discussed below in greater detail, a query base process, such as query base process 10 of FIG. 1, may include one or more nodes communicating with at least one of the multiple data sources to access data therefrom. The query base process may include receiving a query instance. The query instance may be processed to generate one or more relational query instances based at least

in part on the query instance, with each of the one or more relational query instances being distributed among one or more nodes to be processed by at least one node of the one or more nodes to extract data from at least one data source of multiple data sources in communication therewith corresponding to the one or more relational query instances. One or more insights may be generated based on the extracted data. The one or more insights generated may be stored in at least one data source of the multiple data sources, to be retrieved therefrom.

[0049] In some example implementations, the instruction sets and subroutines of the query base process 10, which may be stored on storage device, such as storage device 16, coupled to computer 12, may be executed by one or more processors and one or more memory architectures included within computer 12. In some example implementations, storage device 16 may include but is not limited to: a hard disk drive; a flash drive, a tape drive; an optical drive; a RAID array (or other array); a random access memory (RAM); and a read-only memory (ROM).

[0050] In some example implementations, network 14 may be connected to one or more secondary networks (e.g., network 18), examples of which may include but may not be limited to: a local area network; a wide area network; or an intranet, for example.

[0051] In some example implementations, computer 12 may include a data store, such as a database (e.g., relational database, object-oriented database, triplestore database, etc.) and may be located within any suitable memory location, such as storage device 16 coupled to computer 12. In some example implementations, data, metadata, information, etc. described throughout the disclosure may be stored in the data store. In some example implementations, computer 12 may utilize any known database management system such as, but not limited to, DB2, in order to provide multi-user access to one or more databases, such as the above noted relational database. In some example implementations, the data store may also be a custom database, such as, for example, a flat file database or an XML database. In some example implementations, any other form(s) of a data storage structure

and/or organization may also be used. In some example implementations, the query base process 10 may be a component of the data store, a standalone application that interfaces with the above noted data store and/or an applet/application that may be accessed via client applications 22, 24, 26, 28. In some example implementations, the above noted data store may be, in whole or in part, distributed in a cloud computing topology. In this way, computer 12 and storage device 16 may refer to multiple devices, which may also be distributed throughout the network.

[0052] In some example implementations, computer 12 may execute a query base application (e.g., a query base application 20) examples of which may include, but may not be limited to, e.g., an enterprise query base application, natural language understanding application, voice processing system application, and the like. In some example implementations, the query base process 10 may be accessed via one or more of client applications 22, 24, 26, 28. In some example implementations, query base process 10 may be a standalone application, or may be an applet / application / script / extension that may interact with and/or be executed within the query base application 20, a component of the query base application 20, and/or one or more of client applications 22, 24, 26, 28. In some example implementations, the query base application 20 may be a standalone application, or may be an applet / application / script / extension that may interact with and/or be executed within the query base process 10, a component of the query base process 10, and/or one or more of client applications 22, 24, 26, 28. In some example implementations, one or more of client applications 22, 24, 26, 28 may be a standalone application, or may be an applet / application / script / extension that may interact with and/or be executed within and/or be a component of the query base process 10 and/or the query base application 20. Examples of client applications 22, 24, 26, 28 may include, but may not be limited to, a query base application such as an enterprise query base application, a standard and/or mobile web browser, an email application (e.g., an email client application), a textual and/or a graphical user interface, a customized web browser, a plugin, an

Application Programming Interface (API), or a custom application. The instruction sets and subroutines of client applications 22, 24, 26, 28, which may be stored on storage devices 30, 32, 34, 36, coupled to client electronic devices 38, 40, 42, 44, may be executed by one or more processors and one or more memory architectures incorporated into client electronic devices 38, 40, 42, 44.

[0053] In some example implementations, one or more of storage devices 30, 32, 34, 36, may include but may not be limited to: hard disk drives; flash drives; tape drives; optical drives; RAID arrays; random access memories (RAM); and read-only memories (ROM). Examples of client electronic devices 38, 40, 42, 44 (and/or computer 12) may include, but may not be limited to, a personal computer (e.g., client electronic device 38), a laptop computer (e.g., client electronic device 40), a smart/data-enabled, cellular phone (e.g., client electronic device 42), a notebook computer (e.g., client electronic device 44), a tablet, a server, a television, a smart television, a media (e.g., video, photo, etc.) capturing device, and a dedicated network device. Client electronic devices 38, 40, 42, 44 may each execute an operating system, examples of which may include but may not be limited to, Android™, Apple® iOS®, Mac® OS X®, Red Hat® Linux®, or a custom operating system.

[0054] In some example implementations, one or more of client applications 22, 24, 26, 28 may be configured to effectuate some or all of the functionality of query base process 10 (and vice versa). Accordingly, in some example implementations, the query base process 10 may be a purely server-side application, a purely client-side application, or a hybrid server-side / client-side application that may be cooperatively executed by one or more of client applications 22, 24, 26, 28 and/or the query base process 10.

[0055] In some example implementations, one or more of client applications 22, 24, 26, 28 may be configured to effectuate some or all of the functionality of the query base application 20 (and vice versa). Accordingly, in some example implementations, the query base application 20 may be a purely server-side application, a purely client-side

application, or a hybrid server-side / client-side application that may be cooperatively executed by one or more of client applications 22, 24, 26, 28 and/or the query base application 20. As one or more of client applications 22, 24, 26, 28, the query base process 10, and the query base application 20, taken singly or in any combination, may effectuate some or all of the same functionality, any description of effectuating such functionality via one or more of client applications 22, 24, 26, 28, the query base process 10, the query base application 20, or combination thereof, and any described interaction(s) between one or more of client applications 22, 24, 26, 28, the query base process 10, the query base application 20, or combination thereof to effectuate such functionality, should be taken as an example only and not to limit the scope of the disclosure.

[0056] In example some implementations, one or more of users 46, 48, 50, 52 may access computer 12 and query base process 10 (e.g., using one or more of client electronic devices 38, 40, 42, 44) directly through network 14 or through secondary network 18. Further, computer 12 may be connected to network 14 through secondary network 18, as shown with phantom link line 54. The system of the query base process 10 may include one or more user interfaces, such as browsers and textual or graphical user interfaces, through which users 46, 48, 50, 52 may access the query base process 10.

[0057] In some example implementations, the various client electronic devices may be directly or indirectly coupled to network 14 (or network 18). For example, client electronic device 38 is shown directly coupled to network 14 via a hardwired network connection. Further, client electronic device 44 is shown directly coupled to network 18 via a hardwired network connection. Client electronic device 40 is shown wirelessly coupled to network 14 via wireless communication channel 56 established between client electronic device 40 and wireless access point (e.g., WAP) 58, which is shown directly coupled to network 14. WAP 58 may be, for example, an IEEE 802.11a, 802.11b, 802.11g, Wi-Fi®, RFID, and/or Bluetooth™ (including Bluetooth™ Low Energy) device that may be capable of establishing wireless communication channel 56 between client electronic device 40 and

WAP 58. Client electronic device 42 is shown wirelessly coupled to network 14 via wireless communication channel 60 established between client electronic device 42 and cellular network / bridge 62, which is shown directly coupled to network 14.

[0058] In some example implementations, some or all of the IEEE 802.11x specifications may use Ethernet protocol and carrier sense multiple access with collision avoidance (i.e., CSMA/CA) for path sharing. The various 802.11x specifications may use phase-shift keying (i.e., PSK) modulation or complementary code keying (i.e., CCK) modulation, for example. Bluetooth™ (including Bluetooth™ Low Energy) is a telecommunications industry specification that may allow, e.g., mobile phones, computers, smart phones, and other electronic devices to be interconnected using a short-range wireless connection. Other forms of interconnection (e.g., Near Field Communication (NFC)) may also be used.

[0059] Referring also to the example implementation of FIG. 2, there is shown a diagrammatic view of client electronic device 38. While client electronic device 38 is shown in this figure, this is for example purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible. Additionally, any computing device capable of executing, in whole or in part, query base process 10 may be substituted for client electronic device 38 (in whole or in part) within FIG. 2, examples of which may include but are not limited to computer 12 and/or one or more of client electronic devices 38, 40, 42, 44.

[0060] In some example implementations, client electronic device 38 may include a processor and/or microprocessor (e.g., microprocessor 200) configured to, e.g., process data and execute the above-noted code / instruction sets and subroutines. Microprocessor 200 may be coupled via a storage adaptor to the above-noted storage device(s) (e.g., storage device 30). An I/O controller (e.g., I/O controller 202) may be configured to couple microprocessor 200 with various devices, such as keyboard 206, pointing/selecting device (e.g., touchpad, touchscreen, mouse 208, etc.), custom device (e.g., device 215), USB ports,

and printer ports. A display adaptor (e.g., display adaptor 210) may be configured to couple display 212 (e.g., touchscreen monitor(s), plasma, CRT, or LCD monitor(s), etc.) with microprocessor 200, while network controller/adaptor 214 (e.g., an Ethernet adaptor) may be configured to couple microprocessor 200 to the above-noted network 14 (e.g., the Internet or a local area network).

[0061] As will be discussed below, in some example implementations, the query base process 10 may be integrated into a practical application to at least help, e.g., improve existing technological processes associated with, e.g., querying of multiple data sources necessarily rooted in computer technology.

[0062] It will be appreciated that the computer processes described throughout are not considered to be well-understood, routine, and conventional functions.

The Query Base System:

[0063] The example implementation of FIG. 3 shows a schematic of a network environment (e.g., network environment 300), in accordance with one or more example embodiments of the disclosure. The network environment 300 may implement the query base system 302 (hereinafter, sometimes, simply referred to as “network system” or “system”). In some example implementations, the network environment 300 may be an enterprise network environment. An enterprise network may be generally defined as an enterprise's communications backbone that helps connect computers and related devices across departments and workgroup networks, facilitating insight and data accessibility. The network environment 300 may reduce communication protocols, facilitating system and device interoperability, as well as improve internal and external enterprise data management. The network environment 300 may include local and wide area networks (LAN/WAN), depending on operational and departmental requirements. The network environment 300 may effectively combine and use different device and system communication protocols. For example, the network environment 300 may integrate all systems, including Windows™ and Apple™ computers and operating systems (OS),

Unix™ systems, mainframes and related devices like smartphones and tablets.

[0064] In the example of FIG. 3, the network environment 300 may be a hybrid cloud network with multiple mesh nodes 330, 332, 334 which may be communicating with each other using, for example, Internet 304 as a backbone. Each of the multiple mesh nodes 330, 332, 334 may form a sub-net in the network environment 300. In the shown example, the network environment 300 may include three sub-nets, namely an application sub-net 306, a local sub-net 308, and a cloud sub-net 310. Each of the sub-nets 306, 308, 310 may include respective one or more data sources. In the shown example, the application sub-net 306 may include a data source 312, the local sub-net 308 may include data sources 314, 316, 318 and the cloud sub-net 310 may include data sources 320, 322. It may be appreciated that the network environment 300 may be deployed by an enterprise, for example an organization or a company, which may own or have access to servers located on company property at on-premises location and/or servers in the cloud. The on-premises servers may include some dedicated servers that are physical machines and other dedicated servers that run a virtualization host (vHost or hypervisor) software, such as VMWare or Xen, originally developed by the University of Cambridge Computer Laboratory. The virtualization host software may run several virtual-machine nodes, VM nodes, which may each run applications to service client requests.

[0065] In the example embodiment, the network system 302 may provide enterprise query base (EQB) nodes 330, 332, 334 for each of the sub-nets 306, 308, 310, respectively. For example, the application sub-net 306 may include an application EQB node 330, the local sub-net 308 may include a local EQB node 332, and the cloud sub-net 310 may include a cloud EQB node 334. The application EQB node 330 may be deployed locally with the data source 312 (e.g., locally deployed data source), the local EQB node 332 may be deployed on-premises with the data sources 314, 316, 318 (e.g., on-premises deployed data sources), and the cloud EQB node 332 may be deployed on cloud with the data sources 320, 322 (e.g., cloud deployed data sources). Also, as shown, the application sub-net 306

may include a client application 336 for receiving/generating a query (e.g., variety types of data search queries such as navigational queries, informational queries, and/or transactional queries) for the application EQB node 330. For example, the query may be received as a natural language query such as “show me the sales for our products last year” (e.g., from user speech). Further, the application sub-net 306 may include a graph database 338 for providing metadata for processing the query (e.g., a semantic service may be used with a graph database for metadata or a graph database that stores metadata from the semantic service – the semantic service may be part of an NLP system). As shown, the application EQB node 330 may communicate directly with the graph database 338 and the client application 336 in processing the query. Further, the client application 336 may communicate with the graph database 338 in processing the query.

[0066] In some example implementations, each of the nodes 330, 332 and 334 may interact with respective data source(s) only. That is, the application EQB node 330 may interact with the data source 312 (e.g., containing application data), the local node 332 may interact with the data sources 314, 316, 318 (e.g., containing premises data), and the cloud node 334 may interact with the data sources 320, 322. These nodes 330, 332, 334 may have specific unique permissions allowing each node 330, 332, 334 to interact with its respective data source locations (e.g., application data, premises data, or cloud data). For instance, the cloud node 322 may not be able to interact with premises data as the cloud node 322 may not have permissions to access the on-premises data sources 314, 316, 318. In summary, the query may have a direction sent from a front end to data sources where each EQB node may be connected to data sources within the same sub-network as described in the disclosure (e.g., based or dependent on EQB node’s permissions).

[0067] The system 302 may be provided with multiple nodes, e.g., nodes 330, 332, 334 in different locations such that each of the nodes 330, 332, 334 may become an interface with respective different data source(s). The system 302 may allow for mapping of communication across different parts of the network environment 300, specifically

different data sources. Each of the nodes 330, 332 and 334 may be configured slightly different such that each node may be deployed with its respective location (e.g., based on legal, security policies, etc. relating to the respective location). Each EQB node 330, 332, 334 may provide multiple functionalities, including, but may not be limited to, semantic discover and translation based on rules, in-memory calculation, data source matching, query distribution, query plan construction, and query result integration/caching, etc.

[0068] Example FIG. 4 is a schematic diagram of an example query base system 400 (e.g., enterprise query base system). As shown, the query base system 400 may include a query base subsystem 402 (e.g., enterprise query base subsystem such as an enterprise query base platform) that may interact with external applications and data sources. As shown, the query base subsystem 402 may be in communication with multiple data sources (e.g., data sources 404). As shown, the EQB subsystem 402 may include (at least) two nodes: one node may provide client functionality (referred to as a master client node 406), and another node may provide server functionality (referred to as a server node 408). The two nodes 406 and 408 may interact by means of a network infrastructure, such as via the Internet 304. The data sources 404 may be separated from the master client node 406 via a firewall, and the server node 408 may be located behind the firewall to be able to access and communicate with the data sources 404. In example embodiments where there are at least two nodes (e.g., master client node 406 and server node 408), the master client node 406 may not have a direct connection to the data sources 404 behind the firewall where the server node (e.g., server node 408) may instead be used to interact with the data sources 404 directly. This is important to have separation (e.g., master client node connected to data sources via at least one server node) for security purposes as the data sources (e.g., databases) behind the firewall may not be able to communicate with user(s) directly outside the firewall. This is also important for performance reasons to prevent these data sources from negatively impacting performance (e.g., slow down effect) of the user device and/or the master client node as the server node may be deployed separately with the data source

of the organization/enterprise (e.g., user accesses data via master client node which may indirectly access the data via the at least one server node). Fetching large data from remote data sources over the Internet and processing such data at the master client node may impact performance. Instead, the master client node (and therefore the user device) may interact with data of a query result (as obtained by the server node). In summary, this master/client server node example may be an elastic enabling deployment that may address security and/or performance concerns. In other examples, such as single EQB node examples described in the disclosure, data sources may be directly accessible by the single master EQB node (e.g., one EQB node deployed). In the single node examples, there may be a direct connection to the data sources 404 behind the firewall.

[0069] The query base system 400 may also include an application client 410 to receive a query instance such as from a user (also, sometimes referred to as a user query). The query base system 400 may further include a graph database 412 that may contain metadata (e.g., ontology metadata for semantic functionality) for processing the user query. The graph database 412 may be used for a semantic layer to interpret the query instance (e.g., natural language query instance). The master client node 406 may be configured to interact with the application client 410 as well as with the graph database 412. For example, the application client 410 may receive the query instance and may pass the query instance to the client node 406 in the EQB subsystem 402. The master client node 406 may implement a role mapping/rule mapping logic 414 for processing the received query instance, for example, for not only structured query language (no-SQL) or structured query language (SQL) translation. SQL may be used with databases (e.g., relational databases (RDBMS)) such as PostgreSQL, MySQL, Oracle, Microsoft SQL, and the like whereas no-SQL may be used with databases (e.g., non-relational or distributed databases) such as Redis, RavenDB, Cassandra, MongoDB, BigTable, HBase, Neo4j, CouchDB, and the like. In particular, the master client node 406 may process the query instance to generate one or more relational query instances based at least in part on the query instance, with each of

the one or more relational query instances to be processed by the server node 408. The query instance received may be in the form of a natural language query (NLQ). The query instance may be processed to generate the relational query instances that may be in the form of SQL queries. The master client node 406 may process (and translate) the query instance (e.g., user query as NLQ) to generate relational query instances that may be used for accessing and querying e.g., relational databases. For example, the NLQ query instance may be “show me the sales for my company” and the corresponding SQL relational query instances may be generated as “select sum(invoice.amount) from sales_table” (e.g., NLQ query instance may be translated to the generated SQL relational query instance). The SQL relational query instance may be in different forms such as “select” SQL text format, internal abstract tree structure, and the like. In some examples, the master client node 406 may receive a natural language query request related to obtaining data stored in the multiple data sources 404. In such case, the master client node 406 may parse the natural language query request to generate the query instance. In some example implementations, the master client node 406 may utilize metadata information, from the graph database 412, for parsing of the natural language query (e.g., metadata information may be used to understand and translate natural language query). The graph metadata, as received from the graph database 412, may help the master client node 406 to translate the query using a natural language processing (NLP) service to further disambiguate the query (e.g., NLP may be used for determining intents and disambiguate natural language query). As described in the disclosure, the graph metadata may be ontology metadata that may be in the form of ontology graphs for semantic reasoning (e.g., ontology metadata may include saved ontology graphs used for semantic reasoning). The master client node 406 may use this ontology metadata to parse (e.g., identify vocabulary, meaning, intent, dimension – such as what does “sales” mean based on ontology metadata), understand the NLQ query instance (e.g., may resolve ambiguities in the NLQ query instance), and then mapping in order to generate accurate corresponding relational query instances (e.g., SQL relational query

instances). For example, the query instance may be initially received from the user in natural language (e.g., as a voice or text). Then, the system may consult the graph metadata (e.g., of the graph database 412) to provide NLP and may use mapping to translate the query instance to the relational query instance (e.g., “select” statement). The query instance may be translated to the relational query instance based on the graph metadata (e.g., based on mapping).

[0070] Further, the server node 408 may be configured to optimize the one or more relational query instances to distribute among one or more of the data sources 404 in communication therewith (e.g., optimize may mean use of statistics, histogram, and/or rules to find an optimal plan for distribution and execution of the relational query instances). The server node 408 may include proxy tables 416 and local cache tables 418. The proxy tables 416 (e.g., proxy tables/views) and the local cache tables 418 may include information to build a query plan for distribution of the relational query instances among one or more of the data sources 404. In the examples, the proxy tables 416 may include a partial or a complete map of the data structure in the one or more of the multiple data sources 404. The partial or complete map may be pre-configured based on the data sources 404. For example, the partial or complete map may be setup at a configuration stage or registration stage (e.g., register data sources 404 or specifically register tables of the data sources 404 which may be used to create pre-configured partial or complete map of all data sources 404). The partial or complete map may be updated each time new data is added to data sources 404 or new data sources are added to a collective group of data sources. The partial or complete map is important as the data is not stored in the server nodes. The partial or complete map may be a schema or location mapping to the data structure (e.g., mapping to a table or tables). The partial or complete map (also may be referred to as mapping) assists the server node 408 with generating subqueries that may be sent to data sources 404 or another server node. The results from different data sources 404 may be aggregated at later stages e.g., by the master client node 406.

[0071] The local cache tables 418 may be a data repository which temporarily stores data extracted from the data sources 404 based on the user query for further use such as for another similar query or the like. For example, the local cache tables 418 may be cache tables designed for improving performance. A frequently used data table may be stored as cache tables for improving performance of queries. Generally, it may be understood that the user query may require information from multiple data sources to be processed. The server node 408 may process the relational query instances (corresponding to the user query) to determine from which of the data sources may there be information needed to be gathered for processing of the user query instance. For example, the server node 408 may use the partial or complete map to determine that the locations of the information (e.g., needed for processing the query instance) may be in relational tables t1 and t2. The information may be specifically in column c1 of table t1 and column c1 of table t2. Using the partial or complete map, the server node 408 may determine where table t1 and table t2 reside so that it may generate and send sub-queries to each data source for the relevant data and then join the relevant data (e.g., different sub-query per data source and then join data of same data type). In summary, this example describes a join between two tables (e.g., t1 and t2) from two different data sources (e.g., data source 1 and data source 2, respectively) because the relevant data of these tables (based on query instance) may be joined to create full result (e.g., where part of needed result is in table t1 and part of needed result is in table t2). For example, a sub-query related to table t1 may be sent to data source 1 and another sub-query related table t2 may be sent to data source 2. The syntax for this SQL relational query instance may be “select t1.c1, t2, c1 from t1, t2 where t1.c2 = t2.c2”, “t1.c2 = t2.c2” which means that column c2 from tables t1 and column c2 from table t2 have the same data type. This process of joining tables (e.g., two tables) may occur where the tables have the same data type column.

[0072] It is to be understood that the master client node 406 may communicate with the server node 408 using query language (e.g., translated query language). On the other

hand, the server node 408 may be positioned behind the firewall and may receive the processed query from the master client node 406. The server node 408 may then distribute and interact with the data sources 404. The server node 408 may directly interact with the data sources 404 which may include, for example, but may not be limited to, a file system, and further with an online analytical processing (OLAP) wrapper (which may interact with an OLAP), online transaction processing (OLTP) wrapper (which may interact with an OLTP), No-SQL wrapper (which may interact with a NO-SQL), and a generic open database connectivity (ODBC) connector/wrapper (which may interact with data source 1, data source 2, etc.). Further, the functionalities of the EQB subsystem 402, including the master client node 406 and the server node 408, may be explained in reference to FIG. 5A in the disclosure.

[0073] Example FIG. 5A is a schematic diagram of the EQB nodes 406, 408 of the query base subsystem 402 showing its various sub-components of the query base system 400. FIG. 5A also depicts internal data flow for the EQB subsystem 402. The EQB subsystem 402 is shown to be in external communication with the application client 410, a data processor 504, the graph database 412, data sources 404 and a next EQB node 502. For simplification, the data sources 404 may be considered as regular SQL data sources. As shown, the EQB subsystem 402 may include the master client node 406 and the server node 408. It is to be understood that the EQB subsystem 402 may function either as a standalone client node or a standalone server node if needed, depending on the required role for processing the received query. It may be appreciated that at least one of the multiple data sources may be at least one of a database, a server node, and/or another EQB node.

[0074] As discussed, the network implementing the EQB node (such as, the EQB node 406, 408) may be a directional graph or tree-like mesh that passes the query from the master client node 406 to the data sources 404 through the server node 408. A root of this network tree may generate the query that may be translated and distributed to the leaf nodes (e.g., location of data source) where there may be a clear direction and client/server relationship

(e.g., clear direction from client node which receives natural language, translates natural language, and then distributes relational query instance (or subquery) to different server nodes). In some examples, multiple server nodes 408 may need to be traversed to reach a back end data source having the relevant data/information for processing the query.

[0075] In example embodiments of the disclosure, the EQB subsystem 402 at the root of the network, e.g., the EQB subsystem which may be connected to the application client 410, may include the master client node 406 (and/or act as a client node) to process the query received from the application client 410. The same EQB subsystem 402 may also include the server node 408 (and/or act as a server node) which is in communication with one or more data sources (such as the data sources 404) to extract data for processing of the query received at the master client node 406. In some examples, the EQB subsystem 402 may act as a parent node or set of nodes for the network, which in turn may be connected to one or more other EQB nodes (such as, the next EQB node 502) which may act as sub-nodes for the query base system 400. The one or more enterprise query nodes (such as, the EQB nodes 406, 408 of the EQB subsystem 402) may be arranged in a hierarchical structure (or otherwise). In some example embodiments, the one or more enterprise query nodes (such as, the EQB nodes 406, 408, 502) may be arranged as multiple enterprise query nodes across at least two or more different enterprise networks. These sub-nodes may include functionalities of the server node 408 (and/or act as the server node) and may be disposed in communication with respective one or more data sources to extract relevant data therefrom for processing the query. For example, when the server node 408 communicates with the next EQB node 502, the server node 408 may change its role from working as a “server node” to working as a “client node” and the next EQB node 502 may work as a “server node” such that the server node 408 may act as a client to the next EQB node 502. In examples, the next EQB node 502 may have same components as server node 408. In some examples, the next EQB node 502 may then communicate and work as the “client” with respect to other server nodes until the data is obtained from an appropriate

data source based on the query. In examples, the server node 408 may communicate with the next EQB node 502 and the next EQB node 502 may communicate with other EQB nodes via the Internet 304 directing the relational query instances towards relevant data sources 404 (e.g., via the next EQB node 502 and possibly other EQB nodes). The server node 408 may determine where to direct the relational query instances based on the data source mapping (e.g., using a partial or complete map for determining most efficient pathway to data sources 404). As described in the disclosure, the EQB subsystem 402 may be preconfigured with a pre-determined partial or complete map. The mapping may be done when the EQB subsystem 402 is deployed and may be part of configuration for the EQB subsystem 402 such that all information in the data sources 404 may be mapped prior to use of the EQB subsystem 402. The mapping may be updated/configured at different stages in the future depending on changes to the data sources 404.

[0076] As shown in FIG. 5A, the EQB subsystem 402, or specifically the master client node 406 therein, may receive a query as input from the application client 410. It may be appreciated that the application client 410 (also generally referred to as a client electronic device) may be a user device, such as a personal computer, a smartphone, a tablet computer, or the like which may, optionally, provide a graphical user interface (GUI) to allow a user to input the query therein. Generally, the received query may be a natural language query. In the EQB subsystem 402, the query (e.g., user query), in the form of input (e.g., external signal), may be received by an application query parser 516 in the master client node 406. In some example embodiments, the master client node 406 may be configured to receive the natural language query request related to obtaining data stored in the multiple data sources 404, and parse the natural language query request to generate the query instance. In the example, the generated query instance may have the required syntax to be further converted into an SQL query (as discussed in the disclosure) for processing thereof.

[0077] In particular, as shown in FIG. 5A, the application query parser 516 may convert the user query to an Abstract Semantic Graph (ASG) 518. In particular, the master client

node 406, or the application query parser 516 of the master client node 406, may be configured to generate at least one of an Abstract Semantic Graph (ASG) and/or an Abstract Semantic Tree (AST) for the received natural language query. An ASG, also known as term graph, may be a form of abstract syntax in which an expression of a formal or programming language may be represented by a graph. ASG is typically at a higher level of abstraction than AST, which may be used to express the syntactic structure of an expression or program. ASG 518 may consult with the graph database 412 (e.g., graph database storage for client node), in communication therewith, to get metadata needed for the user query. The generated ASG and/or AST may be implemented to receive the corresponding metadata information for the graph database 412 (e.g., metadata received may be used for semantic translation). Further, the master client node 406 may be configured to utilize metadata information, as received from the graph database 412, for parsing of the natural language query to generate the relational query instance (e.g., metadata may be used to understand and translate natural language query).

[0078] The metadata from the graph database 412 may be stored and combined with heuristic rules for role-based translation. Role-based translation may allow for users with different roles to access different information. For example, role-based translation may be used to understand and determine that “sales” for a CEO may mean the whole company sales while “sales” for a manager may mean for a specific designated group only or specific region (e.g., where manager works with sales in the US only). Role-based translation may include translation of the query instance to generate relational query instance(s) based on user’s role (e.g., role of user requesting query). As described in the disclosure, determining the role-based translation and generating the relational query instances may be based on the metadata and heuristic rules. For example, the relational query instance (e.g., SQL relational query instance) may have different restrictions based on security access (e.g., access to certain data may vary from a CEO to an engineering manager to a secretary for a marketing group) and possibly different semantics (e.g., from one user to another user

based on their role such as meaning of “sales” for sales manager of US division may vary from meaning of sales for a CFO). Further, NLP techniques may be integrated to make translation more robust and elastic. The system may utilize a set of graph database style token languages for interfacing of the NLP unit. The Abstract Semantic Graph (ASG) and the Abstract Semantic Tree (AST) may be used as the internal language for pass through and processing. In some examples, the master client node 406 may receive the Abstract Semantic Graph (ASG) query from the application client 410 (e.g., in form of JavaScript™ Object Notation (JSON) format as user query), may translate the query instance into a relational query instance (e.g., translate ASG into an SQL query by consulting a graph database), and may send the query to the server node 408 (or the query may be sent to data source directly).

[0079] The master client node 406 may be further configured to process (e.g., parse, translate or compile, normalize, etc.) the query instance to generate one or more relational query instances based at least in part on the query instance. The parsing (e.g., may parse different languages such as XML, SQL, language, etc.), translating (e.g., may translate same different languages that were parsed), normalizing or compiling (e.g., normalize or compile to system to be processed as described in the disclosure), and the like may be standard processes but may use a semantic layer (e.g., semantic layer used during translation). Each of the one or more relational query instances may be processed by at least one server node 408. The master client node 406 may utilize role mapping/rule mapping logic 414. Role mapping/rule mapping logic 414 may be used to control the different queries generated based on different roles for context/permission. In the example, the query instance along with the acquired metadata (e.g., metadata received by ASG 518 from graph database 412) may then be passed to an SQL query generator 522 which may make use of the configured role mapping/rule mapping logic 414 (e.g., stored logic) to generate an SQL query (e.g., relational query instance). The generated SQL query may be sent to a query source connector 524. The query source connector 524 may include a map

of the server nodes, and may send the generated SQL query to the relevant server node(s) (such as, the server node 408) which may process the SQL query to extract data from one or more data sources 404 or further pass down the query to the next EQB node 502, as required. That is, the one or more relational query instances may be distributed among one or more server nodes 408 to extract data from the at least one data source (such as, the data sources 404) in communication therewith corresponding to the respective relational query instance received thereby.

[0080] The server node 408 may include an SQL parser 526 that receives a signal, e.g., the SQL query, from the query source connector 524 of the master client node 406. The server node 408 may be configured to optimize the one or more relational query instances to distribute among one or more of the at least one data source in communication therewith and/or at least one of other server nodes. The SQL parser 526 may parse the SQL query into a query tree and may send it to a query optimizer 528. The query optimizer 528 may then consult a data source mapping expert 530 (e.g., stored expert) to build a query execution plan by a plan builder 532. The data source mapping expert 530 may include the proxy tables (such as, the proxy tables 416 of FIGS. 4 and 5A) and local cache tables (such as, the local cache tables 418 of FIGS. 4 and 5A). The plan builder 532 in conjunction with the data source mapping expert 530 may use information from the proxy tables 416 and the local cache tables 418 to understand the structure and layout of the data sources and data contained therein to frame sub-queries for each of the relevant data sources and/or for the next EQB nodes 502 (with other data sources in communication therewith) for extracting data therefrom. The plan output, as generated by the plan builder 532, may be passed to a query distributor 534. The query distributor 534 may use logic for distribution in the plan output and may distribute the query (or queries) to data source connectors 536. The data source connectors 536 may make use of stored data source information in the build plan output to send the query to respective data sources 404 and/or to the next EQB node 502. The data source connectors 536 may further receive the extracted data from each

of the queried data sources 404 and the next EQB node 502. The server node 408 may further include a result engine 538 which may be configured for the integration/aggregation/execution of the extracted data from the data sources 404 and/or from the next EQB node 502. The result engine 538 (in the server node 408) may be disposed in communication with a query result constructor 540 (in the master client node 406). It will be appreciated that any of the elements (and/or their associated roles) in any of the figures may be combined and rearranged without departing from the scope of the disclosure.

[0081] In the master client node 406, the query result constructor 540 may receive extracted data from each of the queried one or more server nodes 408. The query result constructor 540 may aggregate the received extracted data. The data or results from each data source may not be enough such that the query instance and the relational query instance may relate to data from multiple data sources. The query result constructor 540 may be used to aggregate the received extracted data when it comes from multiple data sources (e.g., process sub-query results from different data sources to answer user's query instance). In some example embodiments, the query result constructor 540 of the master client node 406 may be further configured to process the aggregated data to generate a response based on the query instance. The query result constructor 540 may accordingly send an external signal, as output. In some example implementations, the output having the aggregated data, as extracted from the queried data sources 404 and the next EQB node 502, may generally be raw data. The output from the query result constructor 540 may be received by the data processor 504 which may process the received aggregated data in the output to be passed to the application client 410 for perusal of the user (e.g., the result data and a template of display format of the data may be sent to a front end or user interface of the application client 410). For example, the user query may be a request related to a pie-chart for organization of last financial year results. In the example, the extracted raw data may be in the form of values or numbers indicating the relevant financial year results. The

data processor 504 may process the extracted numbers to generate the required pie-chart to be sent to the application client 410, where it may be displayed in the GUI thereof as desired by the user. As shown, the data processor 504 may be positioned external to the EQB subsystem 402. In other examples, e.g., where data template is for data visualization, the data processor 504 may be positioned within the EQB subsystem 402 such that the EQB subsystem 402 may be configured to process the raw data in a format to be received by the application client 410. For example, as described in the disclosure, the data processor 504 may process the result of query (from query result constructor 540) to generate information or data for display in the GUI (e.g., generate graph, table, chart of results based on templates for display).

[0082] Example FIG. 5B is a schematic diagram of a query base system 550 similar to the query base system 400 of FIGS. 4 and 5A. In this example, the query base system 550 may include a single EQB node 552 that may provide the functionalities of the EQB nodes 406, 408 together in one node instead of two separate nodes. In this example, as shown in FIG. 5B, the EQB node 552 may include EQB modules 560, 562 that may respectively provide the functionalities of the EQB nodes 406, 408. For example, a master client module 560 may provide similar functionality to the master client node 406 and a server module 562 may provide similar functionality to the server node 408. As shown, these EQB modules 560, 562 may have the same (or similar) sub-components to the EQB nodes 406, 408, respectively. Accordingly, the disclosure descriptions of these sub-components may be applicable to the functionalities of the master client module 560 and server module 562. FIG. 5B may depict internal data flow for the EQB node 552 similar to the internal data flow of the EQB subsystem 402. In this example, the data flow occurs within the EQB node 552 such that information may be communicated directly between the modules 560, 562 without the need for the Internet 304. For example, the EQB node 552 may have access to one or more in-house data sources without needing to use the Internet 304. In other examples, information may be exchanged between the modules 560, 562 via the Internet.

The EQB node 552 is shown to be in external communication with the application client 410, a data processor 504, the graph database 412, data sources 404 and possibly a next EQB node 554 (via the master client module 560 and server module 562). In some examples, the next EQB node 554 may not be included in the query base system 550. In other examples, the EQB node 552 may communicate with the next EQB node 502 (using server module 562) such that the EQB node 552 acts as client to the next EQB node 502. The EQB node 552 may communicate with the next EQB node 502 via the Internet 304 or without using the Internet 304. In examples, the next EQB node 502 may only include a server module that has same components as the server module 562. In some examples, the next EQB node 502 may then act or function as the “client” with other EQB nodes (these EQB nodes may also only include server modules having same components as server module 562) until the data is obtained from an appropriate data source based on the query.

[0083] Referring now to the example FIG. 6, shown is a schematic diagram of an example process 600 in which one or more proxy tables 416 associated with the server node 408 (as described in the disclosure) may be mapped to one or more base tables 604 associated with a remote data source 610 (e.g., the remote data source 610 may include the base tables 604 and/or procedures). The remote data source 610 may be one of several data sources (e.g., part of the data sources 404). One of the proxy tables 416 (such as, proxy table 606) may refer to a remote data source mapped to one of the base tables 604 (such as, the base table 608) or procedure in an external data source. The utilization of the proxy table (such as, the proxy table 606) for building the plan for distribution of the query is described in the disclosure.

[0084] In some example implementations, the data sources may need to be registered. The system may provide a user interface for configuring/registering new data sources. The configuring/registering of new data sources may be needed to access different data sources (e.g., knowledge of an internet protocol (IP) address of the data source and port number may be used by system to connect with each data source). In this regard, the system may

need to have pre-defined data sources that may be understood by the EQB nodes, or specifically the server nodes, therein. For external data sources, the registration may include a connection definition depending on the data source type. An example of creating a remote server for a data source may be using the following Data Definition Language (DDL) statement as follows:

```
REGISTER SERVER 'RemoteSS'  
CLASS 'EQBODBC'  
USING 'rimu:6666';
```

Wherein '*RemoteSS*' is the name of the remote server, '*EQBODBC*' is a keyword indicating that the remote server is EQB and the connection to it is JDBC (Java Database Connectivity) which is a Java API to connect and execute the query with the database, and *rimu:6666* is the computer name and the TCP/IP port number where the remote server is located. It will be appreciated that languages other than Java™ may also be used without departing from the scope of the disclosure.

[0085] The local proxy table (such as, the proxy table 606) that maps to the remote object may enable location transparency of remote data. The database administrator (DBA) may be able to create the proxy table 606 to access any object (including tables, views, and procedures) that the remote database exports as a candidate for a proxy table. In case of a non-relational data source, local proxy tables 416 (such as, the proxy table 606) may not correspond to a remote relational table. In this example case, the mapping from the non-relational data source to the relational model may be done by an adapter. For example, in case of flat files, a list of .csv files with the same structure may map to a single proxy table. In another example, in case of web services, a set of XML complex elements may be mapped to a set of proxy tables 416. In some example implementations, a specific GUI may be needed to customize the mapping at data source registration for each specific use case. The adapter may be a software layer that allows the enterprise query base (EQB) subsystem 402 to connect to a remote data source; for example, a generic ODBC

connection to remote source, a native data source connection protocol, a specific ODBC/JDBC/ODATA gateway for optimized server access, a Representational State Transfer (REST) API based communication and/or a security layer, etc.

[0086] Example FIG. 7 is a simplified schematic diagram of a query base system 700 depicting propagation of a query therein. As shown, the application client 410 may send the query (e.g., natural language query) to an application query parser 704 (which is similar to application query parser 516 as discussed in the disclosure and may provide similar functionality thereof). In some examples, the query may be constructed against a federated table 705. For example, the query may be created using the federated table (e.g., natural language query created against the federated table 705). During a consumption phase, the federated view may be used (e.g., the federated table 705) with the application query parser 704 for building the queries sent for execution to data sources 404 via a query engine 708. The data sources 404 may be transparent to the query engine 708. Each data source 404 may be any external database system or another EQB node which may act as a proxy for a set of data sources behind a firewall, etc. Communication between the query engine 708 and the data sources 404 may be encapsulated into a query engine operator. The query engine 708 may access the remote data sources 404 via the EQB nodes 710 with their corresponding encapsulated proxy tables 416 mapped to the data sources 404, similar to a local table access, by sending sub-queries thereto, which may be in connection with the data sources 404. The query engine 708 may represent functionality of multiple EQB nodes 710 (e.g., master client node 406 and server node 408 and optionally next EQB node 502) that may be used to enable the EQB nodes 710 to generate a query, distribute query, access data based on query, etc. (as described in the disclosure and may provide functionality thereof). The EQB nodes 710 together with the query engine 708 and the application query parser 704 may form an EQB subsystem 712 (like the EQB subsystem 402). The data sources 404 may be definable as taking input parameters (e.g., the mapped data source object may be a function or procedure with the input parameters). This mapping to a

procedure call may be useful to generate a result set dynamically. Further, in the examples, for satisfying the federation mapping workflows of the system, two kinds of additional objects may be added including, e.g., domain tables to be used for specifying domain of values and lookup tables to be used for associating source column values to a given domain of values. In some examples, low level class encapsulation may be used such that remote data source access may be encapsulated into a scan operator, e.g., “RemoteScan”. The scan operator “RemoteScan” may be a leaf node of a query execution plan such that mapping may refer to a remote table object, a remote procedure or query result, or the like. Below is an example and non-limiting process shown in software pseudo-code for “RemoteScan”:

```

Class RemoteScan
{
    RemoteScan (param1);    //This is a method to bind the proxy table
                            // to a remote database object
                            // param1 will pass in the location and port of
                            // the remote object.

    Open();                //This is to open the connection to the remote
                            //object. When the connection is opened,
                            // the connection will keep the connection
                            // until the data access is fully complete.

    Readrow(N);            //This is to read N rows from the remote
                            // source. It may depend on how the local
                            // batch size is configured for processing. N
                            // may be 1 if multiple rows are not kept for
                            // processing at a time.

    Close();               // Close will disconnect and close for the
                            // current scan. The object can be re-opened
                            // for the second connection and read.

```

```

..... // Additional methods could include convert,
// transform, join, and other processing, etc.
}

```

[0087] In some example implementations, for reducing data transfer, several optimization techniques may be used. For instance, query processing may push remotely as much work as possible to minimize data transfer and increase overall system scalability. This may include, but may not be limited to, full pass-through of single source statements, joins, ordering, sub-queries on sources (e.g., with and without tree transformation), filters (e.g., explicit expressions and implied from transitive closure and join clauses), aggregation (e.g., Eager-Lazy-Aggregation), semi-join reduction (e.g., bind join/partial semi join), etc. The main objectives of the query processing in the federation architecture may include, but may not be limited to, translating the NLP processed semantic representation into a relational SQL representation, maximize performance by reducing data transfer and latency, distribute EQB query processing capabilities to remote data, optimize for the application performance, etc. It will be appreciated that any of the figures discussed throughout may be used in any combination without departing from the scope of the disclosure.

[0088] Referring now to the example FIG. 8, shown is a schematic diagram of a query base system 800 for analysis of data from multiple data sources. The system 800 may include multiple data sources 802 (such as, remote data sources). Herein, the data sources 802 may include databases, data pools, data nodes, etc. The system 800 may include an Enterprise Query Broker (EQB) 804. The EQB 804 may connect to different data sources 802 and may provide a common interface for a distributed data system. The system 800 may also include a semantic query manager (SQM) 806 that is a service that may analyze semantics of the query using ontology before the query may be sent to the EQB 804. The SQM 806 may receive both direct answers and insights and decide the view for the front-end answer. The SQM 806 may pass a user query with its semantic group to the EQB 804,

and the SQM 806 may also obtain direct answers with insights from the EQB 804. The system 800 may further include an analytic service 808 (also referred to as “analytic tool,” “analytics tool,” “analytics service,” etc.) which may pull data from the EQB 804. As shown, the system 800 may further include an insight repository 810 and a forecast models repository 812, which may be treated as separate data sources, that may be joined with other data sources.

[0089] The system 800 may be used to automatically find patterns and correlations across the data based on subject matter. In the disclosed system 800, the analytic service 808 may pull data from the EQB 804 through a common data model instead of understanding all different data systems. The data pulled by the analytic service 808 may be explored intensively and some early correlations may be identified during this process. This is like a data mining stage and the interesting pattern(s) found during this stage may be called insight(s). The insight(s) may be generated by some analytics algorithms (e.g., some typical approaches at data discovery stage may use a variety of models as described in the disclosure such as trend detection, trend change, seasonality, fluctuations, forecast models, threshold alerts, unexpected proportions, outlier analysis, pattern detection, systematic pairwise correlation, casual models, intervention models, etc.). That is, the system 800 may perform various analytics over data to find correlations (e.g., interesting patterns which may be performed through different combinations during a data discovery). The insights discovered at this stage may then be saved by the analytic service 808 into the insight repository 810, with a label denoting its related data field(s). For example, in terms of tagging with related data fields, if data analysis finds an interesting correlation between occupancy rates in hotels and the weather, then the insight (which may be saved in the insight repository) may be tagged with the field occupancy rate. Another correlation may include closing times of an attraction (e.g., at an amusement park) and the impact the closing times may have on crowds in restaurants nearby such that restaurant name may be tagged. Further, another correlation may be between a loyalty program membership tier

and a preferred hotel for stay such that the loyalty membership program membership tier and hotel name may be tagged. When a user query related to some of the data fields is being processed in the EQB 804, the related insight(s) may be recalled as the additional information or context to the user's query may be sent to the front end. The EQB 804 may engage with the insight repository 808 to link insights to cached queries. Herein, the entire process may be scheduled regularly and mechanically. Any new insights from the data may be brought up (e.g., pop-up to user) seamlessly.

[0090] In an example implementation, the system 800 may build and utilize the insight repository 808 with the following pipeline: (i) define and create a library of relevant analytics models, for example: (a) trend detection (for instance, up, down), trend change (for instance, from up to down), seasonality (for instance, weekly fluctuations in bookings such as hotel bookings), forecast models (for instance, sales forecast for the next three months); (b) threshold alert (for instance, price exceeds a certain value), unexpected proportions (for instance, dominant market share), outlier analysis (for instance, one customer returns ten times more orders than the second largest returner), more generally pattern detection (for instance, detect customer service call patterns that lead to higher/lower customer satisfaction), and classification (for instance, determine which category a conversation belongs in positive or negative impact on customer satisfaction); (c) systematic pairwise correlation analysis (for instance, sales of ice cream and hot soups are inversely correlated), causal models (for instance, high occupancy at a hotel is causally related to lower customer satisfaction); and (d) intervention models to extract response signals from past interventions (for instance, response to a marketing intervention); (ii) identify and request data through EQB query; (iii) train models on the requested data, resulting in trained or calibrated models; (iv) name, tag, and save trained/calibrated models to a local database (it may be appreciated that naming convention and tags may make the models easily retrievable); and (v) when, at runtime, a user query expresses the need for some particular analytics, tokens from the parsed user query may be matched against stored

models. If there is a match, the model may be retrieved, the most current data may be requested, the model may be applied to the most current data, and the results may be delivered to the EQB 804; and if the model finds an error greater than a certain threshold when applied to the most current data, the model may be retrained.

[0091] An example use case for the insight repository 810 may be time series forecasting such as sales time series. In an example implementation, a user at a large manufacturing company with multiple business units may prefer to get a sales forecast for the next quarter for an Aerospace business unit. A time series forecasting model may be trained on the Aerospace sales data, with the resulting trained model saved as a Pickle file, for example, "*sales.aerospace.forecast.pkl*." The user query tokens may match this Pickle file. The most current sales data may be requested through the EQB 804. The appropriate current sales data model may be retrieved, unpickled, and applied to the most current data and a forecast may be produced for the required time period. The resulting forecast, with confidence ranges, may be sent to the user through the EQB 804. The insight repository 810 may also be used to enrich customer queries. In the previous example implementation, it may be imagined that, in addition to the requested Aerospace sales forecast, another model may be triggered from predefined rules. For example, other business line sales forecasts or an outlier detection model may be triggered and presented, if applicable, to alert the user to some unusual behavior. Another type of model that may be triggered by a user query may be a what-if simulation model in the form of an intervention model. Again, in the previous example implementation, using the Aerospace sales example, one could request what would happen to the sales forecast if customers were offered a discount or special financing terms. If a discount intervention model exists, trained on customer responses to past discounts, this mode may be used to generate an adjusted forecast when applied to the most current data. Alternatively, depending on predefined rules, an intervention may also be run automatically whenever a user is requesting a sales forecast: the query may then be enriched with the intervention forecast and may be presented to the

user if, for example, the sales lift from applying the discount may be expected to exceed a threshold.

[0092] It may be appreciated that the discovered insights may be relevant for the data prepare process. Compared to the insights, a targeted training may be focused on time series data. The time series data may be historical data (e.g., data histogram) which may be pulled (e.g., downloaded) from the EQB 804. Time series training may be applied to this time series data to generate forecast models. Herein, for example, the forecast model may be trained and tuned based on the time series data extracted from the data source(s) 802 via the EQB 804. These models may be stored (e.g., dumped) in the forecast models repository 812. When a user asks for the prediction-related query, the EQB 804 may call the relevant forecast model (e.g., from the forecast models repository 812) as a plugin operator to forecast the future data based on available data. The forecast model itself may be like other known forecast models (as appreciated by one of skill in the art). In the disclosed system 800, the forecast models may be used as an additional data source. The system 800 may call a particular forecast model in real time data, as discussed in the disclosure. Herein, each forecast model may be updated automatically. In particular, the process of training may be continuous and the forecast models (e.g., from the forecast models repository 812) may be updated automatically with new incoming data. For example, the forecast models may be retrained automatically within a specific period of time depending on amount of data being received and/or changed with respect to the forecast models (e.g., relatively significant change in sales daily, weekly, monthly, etc.). In some examples, the process may be able to create or generate new forecast models based on feedback loops (e.g., using a neural network).

[0093] In the disclosed system 800, the analytic service 808 may pull time series data from the EQB 804 and may dump trained forecast models in the forecast model repository 812. The dumped forecast models may be used as external operators for the EQB 804. The SQM 806 may obtain predicted answers from the EQB 804. External operators may mean

that the forecast models set forth the scope of the data query by the EQB 804 with respect to data that may be relevant to these forecast models based on category or topic of forecast models (e.g., such as sales in a particular region or sales of a particular device, like computer sales only).

[0094] Further, entities may be created in a common data model (CDM) for forecast. The CDM may be a data model that is a common understanding of entities in a given domain and the relationships between these entities. The EQB 804 and the SQM 806 may both utilize the CDM. These entities may then be mapped to the analytic service 808. When a user asks for a prediction via natural language, the EQB 804 may receive an object relation query with the aforementioned entities/attributes. The EQB 804 may then send a request with the entities and attributes to the analytic service 808 to obtain a result. Herein, the analytic service 808 may return a result only if it can find in its data store prediction models that have been tagged with the entities that were received in the request. To build prediction models for the analytic service 808, data may be first pulled into the analytic service 808 via the EQB 804 for exploration and analysis. Then, an appropriate model may be built and trained with historical data. This model may be tagged with the appropriate CDM entities/attributes and uploaded to the data store. A separate service may be setup to retrain the models periodically with new data. This service may use the EQB 804 to pull the latest data relevant to the models, retrain the models, and post the updated models to the data store.

[0095] In an example implementation, an entity called “*predictionSales*” may be created in the CDM with attributes such as amount, date, and “*businessSegment*.” This entity may be mapped to a Representational State Transfer (REST) endpoint of the analytic service 808. A natural language understanding (NLU) module may be trained to understand user requests for prediction. The data science team may use the system to produce a Pickle type of file that may contain a forecast model (e.g., prediction model) for the sales in a business segment aerospace. The file may then be uploaded to the data store and tagged

with “predictionSales.amount” and “predictionSales.businessSegment=aerospace”. When the EQB 804 receives an object relation query containing the “*predictionSales*” entity, the EQB 804 may in turn send a request for aerospace sales along with the time period to the analytic service 808 (which may include stacks of models e.g., sales models such as aerospace sales model, overall sales model, component sales model, regional sales model, etc.). The analytic service 808 may then look up and find the model for aerospace sales, and then use that model to get a prediction for the requested time period and return this result. Accordingly, the analytic service 808 may use each forecast model as an additional data source that may be called in real time.

[0096] In the disclosure, the EQB 804 focuses on conversational queries which may require quick responses as well as possible context. In the EQB 804, query cache may be used to keep a query execution plan and a query context with each query type. In this way, query performance may be achieved by a cached optimal plan. Query context may also be linked within the cached query. For this purpose, the SQM 806 may provide a semantic group which may provide the query context. Query context may also come from feedback usage of user(s) (e.g., rank of frequently asked questions). Accordingly, query context from the SQM 806 and the feedback usage of user(s) may be then linked with the query cache. For example, EQB may store a query cache with a list of queries where each entry may have a set of entity names and relationship labels (e.g., “customer”-[“has”]-> “purchase”, where customer and purchase may be entity names and “has” may be the relationship label to a particular “purchase”). When a new question may be asked by the user, entity names and relationship labels may be generated. These entity names and relationship labels may be used, at least partially, for determining the query context (e.g., query context may also be determined from the SQM 806 and feedback usage). The query context may be used for matching with the cached query. An algorithm (e.g., semantic similarity distance algorithm and/or similarity coefficients such as Jaccard Index) may be used to compute a similarity between the determined query context and a likely cached query such that the closest match

may be selected for the cached query.

[0097] In an example implementation, for example, using the Jaccard index, a user may ask “How much revenue was generated from selling shoes”. Additional information may then be produced through the NLU and the SQM 806 which may look like; “*product:shoes*”, “*metric:revenue*”, “*category:footwear*” together with an object relation query generated by the SQM 806, such as “*GET(sales.amount) FILTER(sales.category(“shoes”).*” The EQB 804 may then further parse the query into sets of entity-relationship labels, e.g. “*sales->[“in”]->category*”, to form a query context. In this case, matching query context with cached context with Jaccard Index may result in a Jaccard coefficient of 0.5. Then, the system 800 may try compute the Jaccard coefficient of the query context and entities of other queries in the query cache and then may take the query with the highest coefficient. This may be represented in tabular form, as shown below:

Entities	Query Context
<u>entity: sales</u>	<u>product: shoes</u>
<u>entity: category</u>	<u>metric: revenue</u>
<u>relation: in</u>	<u>category: footwear</u>

Original query text	Parameterized query text	Entities	Context
<u>GET(receipt.amount)</u>	<u>GET(receipt.amount)</u>	<u>entity:receipt</u>	<u>product: socks</u>
<u>FILTER(receipt.item(“socks”))</u>	<u>FILTER(receipt.item(@))</u>	<u>entity:item</u>	<u>metric: revenue</u>
		<u>relation:has</u>	<u>category:footwear</u>

[0098] As discussed, data exploration may lead to the discovery of interesting patterns or relationships within the data. Such discoveries, called insights, may be published to a repository tagged with the relevant CDM entities and attributes. When the EQB 804 receives an object relation query, apart from executing the object relation query, the EQB 804 may also query the analytic service 506 for any insights/additional information that may have been published for the entities in the object relation query. If available, then the

EQB 804 may return results for the requested query as well as the insight as a part of its response. In some example implementation, the data science team may use the system to publish some information tagged with the CDM attribute “*hotelOccupancy*” and “*hotelName*”. When the user requests occupancy rate of a hotel, the EQB 804 may receive the object relation query for this request. The EQB 804 may send a request with the attributes “*hotelOccupancy*” and “*hotelName*” to the analytic service 808. The analytic service 808 may then respond with the additional information which may be a comparison of occupancy rates of all the hotels in the property. The user may then receive the occupancy rate for the requested hotel and a chart showing the occupancy rates for all the hotels in the property.

[0099] Referring now to the example FIG. 9, shown is query cache in a EQB (such as, the EQB 804 of FIG. 8). Herein, the query cache may be in the form of a data table. As shown, the query cache may include a “query entry” which may include query text that parameterize literal values in the query (e.g., filter values of the query). This may be query text such as object relation query text with parameterize literal values. For example, “the sales in June”, June may be a filter which may be parameterized. The query may be changed to “the sales in @”. In this case, “the sales in June” and “sales in May” may share a same query entry in the query cache. The query cache may also include a “text hash number” which may include a hash value (e.g., numeric value of a fixed length that uniquely identifies data) that may be generated based on the query text. The hash value may be used for a quick search for a query in the query cache. The hash value may provide a relatively unique value in terms of probability. The hash value may be generated using a hash function. There are several types of hash functions (e.g., algorithms) available. Some examples of hashing functions may include hashing variable length-data (e.g., query text) such as middle and ends, character folding, word length folding, radix conversion hashing, rolling hash, and the like. These hash functions may be similar in results but may differ on performance versus uniqueness/collision tradeoffs. For instance, in example FIG. 9, the

text hash number may be “*m*”. Query matching may occur based on a relationship of query text with text hash number. The query cache may further include “execution times” which may be a value to denote how many times the query had been asked by the user. For instance, in example FIG. 9, the execution times may be “*n*” which may refer to popularity of the query. The query cache may further include “linked query plan” which may cache the best query plan to run the query (e.g., best way to run and distribute query). For instance, in example FIG. 9, the linked query plan may include 3 *ranked query plans* each with “*Plan ID – Plan text*”. As discussed, SQM (such as, the SQM 806 of FIG. 8) may provide a semantic group for each query in the query cache, which may be the context the query may apply. For instance, in example FIG. 9, this may be referring to a “semantic query type” such as “*semantic category*” (e.g., query grouping), in the query cache. Further, as disclosed, analytic service (such as, the analytic service 808 of FIG. 8) may link the insights and related queries by a related query link for each query entry. These insights may refer to insights produced by data exploration process when data mining using the EQB (as described in the disclosure) and related queries may be the queries provided by the SQM. The analytic service may then link these insights to related queries using “related query link” in the query cache. For instance, in example FIG. 9, the “related query link” may be “*query numbers or link*” (e.g., “Query number – correlation index”).

[00100] Referring now to the example FIG. 10, shown is a schematic diagram of a query workflow 1000 (which in some example implementations may be executed at least in part by, e.g., query base process 10) in EQB in query cache. Herein, the workflow may be as follows. A query parser 1002 may parse an object relation query in a string/text form into a named token tree that may be called an “Atom Tree” which may be an implementation instance of an Abstract Semantic Tree (AST). The Atom Tree may be a n-ary tree data structure where each node may have a name of the entity or attribute extracted from the object relation query. Then, an entity tree builder 1004 may transform the Atom Tree further by retrieving information about how the entities and attributes are related and

mapped to the underlying data. Such information may be retrieved from a configuration database 1006 which may provide a “config” file in JSON format to a configuration builder 1008 to provide the entity configuration information to the entity tree builder 1004. This transformation may result in a new n-ary tree data structure called the “Entity Tree”, which may serve as mapping entity definitions with the underlying data. A Learning Recommendation Engine 1010 may add related insights to the Entity Tree and store a copy in plan cache 1011. The Learning Recommendation Engine 1010 may be where insight(s) and forecast(s) may sit in the EQB. The Entity Tree with related insights may be processed by a query tree builder 1012 to provide “Query Trees”, which in turn may be processed by a query tree group finder 1014 to form “Grouped Query Trees”. Further, a query tree optimizer 1016 may process the Grouped Query Trees to provide “Optimized Query Trees”. A query tree transformer 1018 may transform the Optimized Query Trees to generate “Physical Plans”. The Physical Plans may be fed by a physical plan runner 1020 to an Sql query runner 1022 which generates an “Sql Model” to be fed to an Sql Query builder 1024 which in turn may generate an “Sql Statement” to be fed back to the Sql query runner 1022 to generate relevant “Data”.

[00101] Referring now to the example of FIG. 11, shown is a schematic diagram for an example implementation of an entity model 1100. Given the entity model 1100, the question: *Get the total sales amount where the category type is of “HotelSales”* may be translated into the following object relation query as also shown in FIG. 11:

Object Relation Query:

GET(SUM:Sale.Receipt.Amount) FILTER(Sale.Category.Type(“HotelSales”))

GET is the keyword for the projection

SUM is the keyword for summation aggregation

FILTER is the keyword for applying conditional filters

Atom Tree:

Sale

```

|---
|--- Receipt
|--- Amount
|--- Category
|--- Type("HotelSales")

```

Entity Tree:

Sale (projection: tbl_sales_header)

```

|---
|--- Receipt (join: tbl_sales_line)
|--- Amount (projection: tbl_sales_line.amount)
|--- Category (join: tbl_category)
|--- Type("HotelSales") (filter: tbl_category.type = "HotelSales")

```

[00102] Herein, the semantic side may provide information such as category (e.g., sales) of question from user. The query cache may include this semantic category (e.g., marker) for the query/question. By categorizing the query with a semantic grouping provided by SQM (such as, the SQM 806 of FIG. 8), query base system (such as, the system 800 of FIG. 8) may access only the insights relevant to the semantic category and skip the irrelevant insights. Thus, the query may be linked to relevant insights. Further, the query cache may be created by EQB (such as, the EQB 804 of FIG. 8) which may be distributed to different nodes. In addition to categories, the query cache may also include links to other queries. Also, the entity tree may be used to find related queries based on a selected algorithm (e.g., using semantic distance which may be preferred or using other similarity coefficients such as Jaccard Index to get an index for ranking the related queries from the query cache). As appreciated by one of skill in the art, other algorithms like or similar to semantic distance and/or similar coefficients (e.g., Jaccard Index) may be used for processing and/or for matching. The output may be then processed in addition to the original query and the final results may be presented as recommendations.

The Query Base Process:

[00103] As discussed in the disclosure and referring also at least to the example implementations of Figs. 3-13, a query base process 10 may include each of one or more nodes communicating with at least one of the multiple data sources to access data therefrom. The query base process 10 may further include receiving a query instance. The query instance may be processed to generate one or more relational query instances based at least in part on the query instance, with each of the one or more relational query instances being distributed among one or more nodes to be processed by at least one node of the one or more nodes to extract data from at least one data source of multiple data sources in communication therewith corresponding to the one or more relational query instances. One or more insights may be generated based on the extracted data. The one or more insights generated may be stored in at least one data source of the multiple data sources, to be retrieved therefrom.

[00104] The disclosure further provides the query base process 10 that may be used for analysis of data from multiple data sources. Example FIG. 12 is an example flowchart of the query base process 10 for analysis of data from multiple data sources, according to example embodiments of the disclosure. The query base process 10 may communicate 1202 (e.g., via one or more nodes) with at least one data source of the multiple data sources to access data therefrom. The query base process 10 may receive 1204 a query instance. In some example implementation, the query context for the query instance may be determined. The query context may be determined from the identified entities and attributes therefrom. Further, one or more predefined parameters may be filtered from the query instance. In an example, receiving 1204 the query instance may include receiving 1224 a natural language query request related to obtaining data stored in the multiple data sources. The query base process 10 may process 1206 the query instance to generate one or more relational query instances based at least in part on the query instance, with each of the one or more relational query instances to be processed by at least one node of the one or more

nodes. In one or more examples, processing 1206 the query instance may include parsing 1226 the natural language query request to generate the one or more relational query instances by utilizing metadata information (e.g., metadata information may be used to understand and translate natural language query request) and the query base process 10 may generate 1228 an Abstract Semantic Graph (ASG) and/or an Abstract Semantic Tree (AST) for the natural language query request to receive the metadata information (e.g., metadata information may be used for semantic translation). The query base process 10 may distribute 1208 the one or more relational query instances among one or more nodes to extract data from the at least one data source in communication therewith corresponding to the respective relational query instance received thereby. The query base process 10 may process 1210 one or more relational query instances at one or more nodes to extract data. In one or more examples, after the query base process 10 distributes 1208 the one or more relational query instances, the query base process 10 may optimize 1230 the relational query instances to distribute among one or more of the at least one data source of multiple data sources and at least another node (e.g., another server node). The query base process 10 may process 1210 one or more relational query instances after optimizing 1230 the relational query instances. The query base process 10 may generate 1212 one or more insights based on the extracted data. The generated one or more insights may be stored 1214 in at least one data source of the multiple data sources, to be retrieved therefrom. The generated one or more insights may be further associated to the corresponding query instance received. The query instance received may be cached, with the one or more insights associated therewith, in at least one data source of the multiple data sources, to be retrieved therefrom. The cached query instance may be linked to the query context based on a matching algorithm. The matching algorithm may be a semantic matching algorithm.

[00105] Example FIG. 13 is an example flowchart of the query base process 10 for further re-training of a model for analysis of data from multiple data sources, according to example embodiments of the disclosure. The query base process 10 may identify 1302

entities in the query instance and attributes associated therewith. The query base process 10 may retrieve 1304 historic data from at least one data source of the multiple data sources related to one or more of the entities and attributes identified. At least one model may be trained 1306 based on the historic data retrieved. Tags may be associated 1308 with the at least one trained model indicative of corresponding one or more of the entities and attributes. The at least one trained model may be stored 1310 with associated tags in at least one data source of the multiple data sources, to be retrieved therefrom. The query base process 10 may monitor 1312 the at least one data source of the multiple data sources for updates in the historic data. Monitoring the at least one data source of the multiple data sources further comprises periodically monitoring the at least one data source of the multiple data sources for updates in the historic data. The entities and attributes associated with the updates in the historic data may be determined 1314. The at least one model, with associated tags related to the determined entities and attributes, may be re-trained 1316 based on the updates in the historic data.

[00106] The query base system may be configured to receive a query instance. The enterprise query base may be further configured to process the query instance to generate one or more relational query instances based upon, at least in part, the query instance, with each of the one or more relational query instances to be processed by at least one node of one or more nodes, wherein the one or more relational query instances may be distributed among the one or more nodes to extract data from the at least one data source in communication therewith corresponding to the one or more relational query instances. The enterprise query base may be further configured to receive extracted data from each of the one or more nodes queried. The computing system may further include an analytic service in communication with the enterprise query base. The analytic service may be configured to generate one or more insights based on the extracted data. The analytic service may be further configured to store the one or more insights generated in at least one data source of the multiple data sources, to be retrieved therefrom.

[00107] One or more of the following example features may be included. The enterprise query base may be further configured to identify entities in the query instance and attributes associated therewith. The enterprise query base may be further configured to retrieve historic data from the at least one data source of the multiple data sources related to one or more of the entities and attributes identified; train at least one model based on the historic data retrieved; associate tags with the at least one trained model indicative of corresponding one or more of the entities and attributes; and store the at least one trained model with associated tags in at least one data source of the multiple data sources, to be retrieved therefrom. The enterprise query base may be further configured to monitor the at least one data source of the multiple data sources for updates in the historic data; determine entities and attributes associated with the updates in the historic data; and re-train the at least one model, with associated tags related to the determined entities and attributes, based on the updates in the historic data. The enterprise query base may monitor the at least one data source of the multiple data sources which may further include the enterprise query configured to periodically monitor the at least one data source of the multiple data sources for updates in the historic data. The enterprise query base may be further configured to associate the generated one or more insights to the corresponding query instance received; and cache the query instance received, with the one or more insights associated therewith, in at least one data source of the multiple data sources, to be retrieved therefrom. The enterprise query base may be further configured to determine a query context for the query instance and link the cached query instance to the determined query context based on a matching algorithm. The matching algorithm may be a semantic matching algorithm. The query context may be determined from identified entities in the query instance and attributes associated therewith. The enterprise query base may be further configured to filter one or more predefined parameters from the query instance. The enterprise query base may be further configured to generate a hash value for the query instance received. The enterprise query base may be configured to use the hash value to find a query cache that

matches the hash value. The enterprise query base may be further configured to generate a query plan based on the processing of the query instance. The enterprise query base may be further configured to process another related query instance based on the query plan. The query instance may be a query text. The one or more nodes may be arranged in a hierarchical structure. The one or more relational query instances may be distributed among the one or more nodes directionally from upper nodes to lower nodes. The enterprise query base may be further configured to process the extracted data to generate a response based on the query instance.

[00108] The disclosed query base system may provide a common gateway and port to bridge mined insights and trained models to be parts of data sources of an ad-hoc query system. The query base system has the ability to integrate data from different data sources. However, some of the data may need intensive calculations before the disclosed system may find some interesting patterns. On the other hand, some data needs may work with analytic models for prediction and planning, etc. The query base system may provide the combination or integration of analytics with enterprise query broker (EQB) processing. In contradistinction, the known example solutions of using a common technique may not integrate the analytic results with real time transactional data. The query base system may connect analytic service with data via EQB which may provide caching and linking of the insights to the query, and thus may help build insight repository and prediction models. This may help to integrate real time data with forecast and insights.

[00109] The query base system may be based on a mesh concept, that may work as a gateway for a data source federation, may distribute query, and may aggregate the result from different data sources. In example embodiments, the mesh concept may be employed in a query-base type of system as providing different purposes and applications. Benefits to using mesh concept may relate to infrastructure level and application level. For example, the query base system may provide a query engine that may be used to query different types of databases (e.g., query base system may not store data with itself but may instead connect

to “different” data sources where “different” may mean different locations of the data source and/or different types of data sources). In contrast from other data base systems, the query base system may have only network level of access plan of data but may not have a store level access plan. The query base system may cross different networks even including Internet to access data sources that may be behind a different firewall. The query base system may use a mesh network structure for allowing a new data source to plug into the query base system quickly and easily as well as flexible and elastic. The elastic mesh concept may address various complex enterprise data source issues which may include hybrid data store including on premise and in cloud, and/or systems across geographic regions and different firewalls, etc. The mesh structure or layout may help to solve complicated enterprise scenarios where data may be scattered in different places such as cloud, on premise, or other locations. The system may cross different networks including internet and one or more access data sources behind a different firewall. This capability may provide operational and cost benefits, but may also support the development and deployment of the next generation of analytical applications which may require the ability to access, synthesize, and integrate data from multiple systems in real-time regardless of where the data is located or what systems are generating it (e.g., different server nodes may integrate data residing outside the EQB subsystem or EQB platform). These applications’ ability for cross-functional data in near-real time may drive new demand for data federation capabilities (e.g., accessing integrated data from multiple sources without creating redundant data stores).

[00110] The query base system may use different operators for querying data from different data sources (e.g., direct query base may access data sources transparently by query), which is in contradistinction to some known systems which may simply store data (e.g., some systems may allow for data to be copied or replicated from other data sources). In the system, each EQB node may have multiple data source connectors to connect to different data sources which are accessible by the particular EQB node. If a data source

cannot be accessed by the current EQB node, another new EQB node may be deployed in the system/network accessible to the data source. The previous EQB node may communicate and/or interact with the new EQB node to access the data source (e.g., the new EQB node may represent one or more data sources behind a same firewall). Thus, the system may provide functionality of being able to query all individual data across different data sources (e.g., no need to copy or replicate data). That is, the process may include translating a query, then query data sources (e.g., databases) directly via one or more nodes to obtain result(s).

[00111] Further, the system may integrate with Natural Language Processing (NLP), which may handle semantic mapping and may remove the barrier for the application to query different data sources with different query languages. The query base system may use semantic operations that may be adapted for different types of data sources, for example depending on type of data source, its network/geographic location, available access type, etc. This means that the system may be able to generate a subquery for one data source that may be completely different from a subquery generated for a second different data source, thus enabling the system to function across different types of networks and also different types of data sources. In some examples, the EQB nodes may use a specific protocol depending on data source (e.g., specific protocol may enable EQB node to communicate with a specific type of data source). In the system, the EQB node may be configured to translate and distribute a query into a relational SQL query and seamlessly access data residing in data sources outside thereof by distributing the query to different data sources. The EQB node may offer a translation framework for native expressions. For example, the EQB node may offer a translation framework for mapping a federation function to data source specific expressions taking into consideration, e.g., native dialects and semantic native expression differences. Further, the EQB node may provide a way for pass-through of expressions or execution of a complete statement in a data source's native syntax in order to use database specific native functions/operators and for performance reasons.

[00112] The EQB node may be a query base which may work as a gateway for data source federation, may distribute the query, and may aggregate the result from different data sources. The EQB node may parse the query, from the application client, into ASG and may translate graph database semantics into relational query instance(s) based on role mapping/rule mapping logic. Each EQB subsystem may include the client node and the server node, where each node may run independently depending on the need. The client node of the EQB subsystem may send the relational query instance(s) to the server node (or single data source port). The client node may further format/construct result(s) catered for the application client.

[00113] The system may be one directional in the sense that the query may be pushed down to the data source(s), whereas some example database management systems may be bi-directional that may obtain a query result and may then send the query result back, and if there is still missing information, such example systems may have to send a new query. In contrast, the disclosed system may send out one query that may move down chain towards different EQB nodes that may be able to pass the query in a direction or directions towards multiple sources that may have information needed for responding to the query. The system may send multiple queries down the chain to obtain query results from one source or more than one source. In the system, the query may not actually include the calculated route information but instead, each EQB node may keep the destination data source mapping and may optimize the query plan itself. In the system, each EQB node may be a sub-query engine and may process the sub-query (e.g., filtering, joining, aggregating, etc.).

[00114] The query base system may be a unique query processing mechanism which may translate a natural language processing (NLP) query into data queries and may distribute the query to different data sources intelligently. The system may use a query processing node mesh architecture which makes accessing the data sources elastic and robust. The system may not actually duplicate the data from the data sources but instead

may query the real time data directly via the query pipeline. The system may make use of a query communication mesh concept which may provide a way to access and process integrated data from multiple sources without creating redundant data stores. This may provide a way to distribute and push down the query to the data stored at its native location and avoid moving massive sensitive data across the network. Further, the system may encapsulate the communication of different data sources and may provide a unified interface for data access and interchange. The system may integrate learning optimization processing to the distributed query plan based on heuristic rules as well as query execution metrics, which may tune the query plan automatically. Thus, the system may provide a unified view of data across multiple servers to access data in each server, transparently, without copying.

[00115] In some examples, the query base system may interact with a data lake (e.g., pull and copy all data into a relatively large repository) with data source replication. For example, the query base system may connect data from different data sources to the same data lake (e.g., extract, translate, and load data from different data sources into one centralized data lake). For this example, the process may be similar to what is described in the disclosure except interactions may be with a single data lake (e.g., EQB nodes may interact with data in a data lake). A benefit of this process may be that data may be easily accessed from the data lake.

[00116] In other examples, query cache and query/data encryption may be used. This may improve performance by running previously executed queries found in query cache. Query/data encryption may relate to security where encryption may be added to data protected data during transmission (e.g., prevent chance of data being obtained through attack during transmissions).

[00117] The terminology used herein is for the purpose of describing particular implementations only and is not intended to be limiting of the disclosure. As used herein, the singular forms "a", "an" and "the" are intended to include the plural forms as well,

unless the context clearly indicates otherwise. As used herein, the language “at least one of A, B, and C” (and the like) should be interpreted as covering only A, only B, only C, or any combination of the three, unless the context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising," when used in this specification, specify the presence of stated features, integers, steps (not necessarily in a particular order), operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps (not necessarily in a particular order), operations, elements, components, and/or groups thereof.

[00118] The corresponding structures, materials, acts, and equivalents (e.g., of all means or step plus function elements) that may be in the claims below are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the disclosure has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the disclosure in the form disclosed. Many modifications, variations, substitutions, and any combinations thereof will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the disclosure. The implementation(s) were chosen and described in order to explain the principles of the disclosure and the practical application, and to enable others of ordinary skill in the art to understand the disclosure for various implementation(s) with various modifications and/or any combinations of implementation(s) as are suited to the particular use contemplated.

[00119] Having thus described the disclosure of the present application in detail and by reference to implementation(s) thereof, it will be apparent that modifications, variations, and any combinations of implementation(s) (including any modifications, variations, substitutions, and combinations thereof) are possible without departing from the scope of the disclosure defined in the appended claims.

What Is Claimed Is:

1. A computer-implemented method for analysis of data from multiple data sources, the computer implemented method comprising:

- receiving a query instance;
- processing the query instance to generate one or more relational query instances based upon, at least in part, the query instance;
- distributing the generated one or more relational query instances among one or more nodes;
- processing each of the generated one or more relational query instances by at least one node of one or more nodes to extract data from at least one data source of multiple data sources in communication therewith corresponding to the one or more relational query instances;
- generating one or more insights based on the extracted data; and
- storing the one or more insights generated in relevant cache to be retrieved therefrom.

2. The method according to claim 1 further comprising identifying entities in the query instance and attributes associated therewith.

3. The method according to claim 2 further comprising:

- retrieving historic data from at least one data source of the multiple data sources related to one or more of the entities and attributes identified;
- training at least one model based on the historic data retrieved;
- associating tags with the at least one trained model indicative of corresponding one or more of the entities and attributes; and
- storing the at least one trained model with associated tags in at least one data source of the multiple data sources, to be retrieved therefrom.

4. The method according to claim 3 further comprising:
 - monitoring the at least one data source of the multiple data sources for updates in the historic data;
 - determining entities and attributes associated with the updates in the historic data; and
 - re-training the at least one model, with associated tags related to the determined entities and attributes, based on the updates in the historic data.
5. The method according to claim 4, wherein the monitoring the at least one data source of the multiple data sources further comprises periodically monitoring the at least one data source of the multiple data sources for updates in the historic data.
6. The method according to claim 1 further comprising:
 - associating the generated one or more insights to the corresponding query instance received; and
 - caching the query instance received, with the one or more insights associated therewith, in an enterprise query base (EQB) system, to be retrieved therefrom.
7. The method according to claim 6 further comprising:
 - determining a query context for the query instance; and
 - linking the cached query instance to the query context based on a matching algorithm.
8. The method according to claim 7, wherein the matching algorithm is a semantic matching algorithm.

9. The method according to claim 7, further comprising identifying entities in the query instance and attributes associated therewith, wherein the query context is determined from the identified entities and attributes therefrom.
10. The method according to claim 1 further comprising filtering one or more predefined parameters from the query instance.
11. The method according to claim 1 further comprising:
 - generating a hash value for the query instance received; and
 - using the hash value to find a query cache that matches the hash value.
12. The method according to claim 1 further comprising:
 - generating a query plan based on the processing of the query instance; and
 - processing another related query instance based on the query plan.
13. The method according to claim 1, wherein the query instance is a query text with literal parameters filtered.
14. The method according to claim 1 further comprising processing the extracted data to generate a response based on the query instance.
15. A computing system for analysis of data from multiple data sources, comprising:
 - an enterprise query base configured to communicate with at least one data source of the multiple data sources to access data therefrom, the enterprise query base is configured to:
 - receive a query instance;

process the query instance to generate one or more relational query instances based upon, at least in part, the query instance, with each of the one or more relational query instances to be processed by at least one node of one or more nodes, wherein the one or more relational query instances are distributed among the one or more nodes to extract data from the at least one data source in communication therewith corresponding to the one or more relational query instances; and

receive extracted data from each of the one or more nodes queried;

and

an analytic service in communication with the enterprise query base, wherein the analytic service is configured to:

generate one or more insights based on the extracted data; and

store the one or more insights generated in at least one data source of the multiple data sources, to be retrieved therefrom.

16. The system according to claim 15, wherein the enterprise query base is further configured to identify entities in the query instance and attributes associated therewith.

17. The system according to claim 16, wherein the enterprise query base is further configured to:

retrieve historic data from the at least one data source of the multiple data sources related to one or more of the entities and attributes identified;

train at least one model based on the historic data retrieved;

associate tags with the at least one trained model indicative of corresponding one or more of the entities and attributes; and

store the at least one trained model with associated tags in at least one data

source of the multiple data sources, to be retrieved therefrom.

18. The system according to claim 17, wherein the enterprise query base is further configured to:

monitor the at least one data source of the multiple data sources for updates in the historic data;

determine entities and attributes associated with the updates in the historic data; and

re-train the at least one model, with associated tags related to the determined entities and attributes, based on the updates in the historic data.

19. The system according to claim 18, wherein the enterprise query base monitors the at least one data source of the multiple data sources further comprises the enterprise query base configured to periodically monitor the at least one data source of the multiple data sources for updates in the historic data.

20. The system according to claim 15, wherein the enterprise query base is further configured to:

associate the generated one or more insights to the corresponding query instance received; and

cache the query instance received, with the one or more insights associated therewith, in at least one data source of the multiple data sources, to be retrieved therefrom.

21. The system according to claim 20, wherein the enterprise query base is further configured to determine a query context for the query instance and link the cached query instance to the determined query context based on a matching algorithm.

22. The system according to claim 21, wherein the matching algorithm is a semantic matching algorithm.
23. The system according to claim 21, wherein the query context is determined from identified entities in the query instance and attributes associated therewith.
24. The system according to claim 15, wherein the enterprise query base is further configured to filter one or more predefined parameters from the query instance.
25. The system according to claim 15, wherein the enterprise query base is further configured to generate a hash value for the query instance received, wherein the enterprise query base is configured to use the hash value to find a query cache that most closely matches the hash value.
26. The system according to claim 15, wherein the enterprise query base is further configured to generate a query plan based on the processing of the query instance, and wherein the enterprise query base is further configured to process another related query instance based on the query plan.
27. The system according to claim 15, wherein the query instance is a query text with literal parameters filtered.
28. The system according to claim 15, wherein the one or more nodes are arranged in a hierarchical structure, and wherein the one or more relational query instances are distributed among the one or more nodes directionally from upper nodes to lower nodes.

29. The system according to claim 15, wherein the enterprise query base is further configured to process the extracted data to generate a response based on the query instance.

30. A computer-implemented method for analysis of data from multiple data sources, the computer implemented method comprising:

- receiving a query instance;
- processing the query instance to generate one or more relational query instances based upon, at least in part, the query instance;
- processing each of the generated one or more relational query instances by at least one node of one or more nodes;
- distributing the one or more relational query instances among the one or more nodes to extract data from the at least one data source;
- receiving the extracted data from each of the one or more nodes queried;
- generating one or more insights based on the extracted data; and
- storing the one or more insights generated in at least one data source of the multiple data sources, to be retrieved therefrom.

31. The method according to claim 30 further comprising identifying entities in the query instance and attributes associated therewith.

32. The method according to claim 31 further comprising:

- retrieving historic data from at least one data source of the multiple data sources related to one or more of the entities and attributes identified;
- training at least one model based on the historic data retrieved;
- associating tags with the at least one trained model indicative of corresponding one or more of the entities and attributes; and

storing the at least one trained model with associated tags in at least one data source of the multiple data sources, to be retrieved therefrom.

33. The method according to claim 32 further comprising:

monitoring the at least one data source of the multiple data sources for updates in the historic data;

determining entities and attributes associated with the updates in the historic data; and

re-training the at least one model, with associated tags related to the determined entities and attributes, based on the updates in the historic data.

34. The method according to claim 30 further comprising:

associating the generated one or more insights to the corresponding query instance received; and

caching the query instance received, with the one or more insights associated therewith, in an enterprise query base (EQB) system, to be retrieved therefrom.

35. A computer program product residing on a computer readable storage medium having a plurality of instructions stored thereon which, when executed across one or more processors, causes at least a portion of the one or more processors to perform operations comprising:

receiving a query instance;

processing the query instance to generate one or more relational query instances based upon, at least in part, the query instance;

processing each of the generated one or more relational query instances by at least one node of one or more nodes;

distributing the one or more relational query instances among the one or more nodes to extract data from the at least one data source;
receiving the extracted data from each of the one or more nodes queried;
generating one or more insights based on the extracted data; and
storing the one or more insights generated in at least one data source of multiple data sources, to be retrieved therefrom.

36. The computer program product according to claim 35 further comprising identifying entities in the query instance and attributes associated therewith.