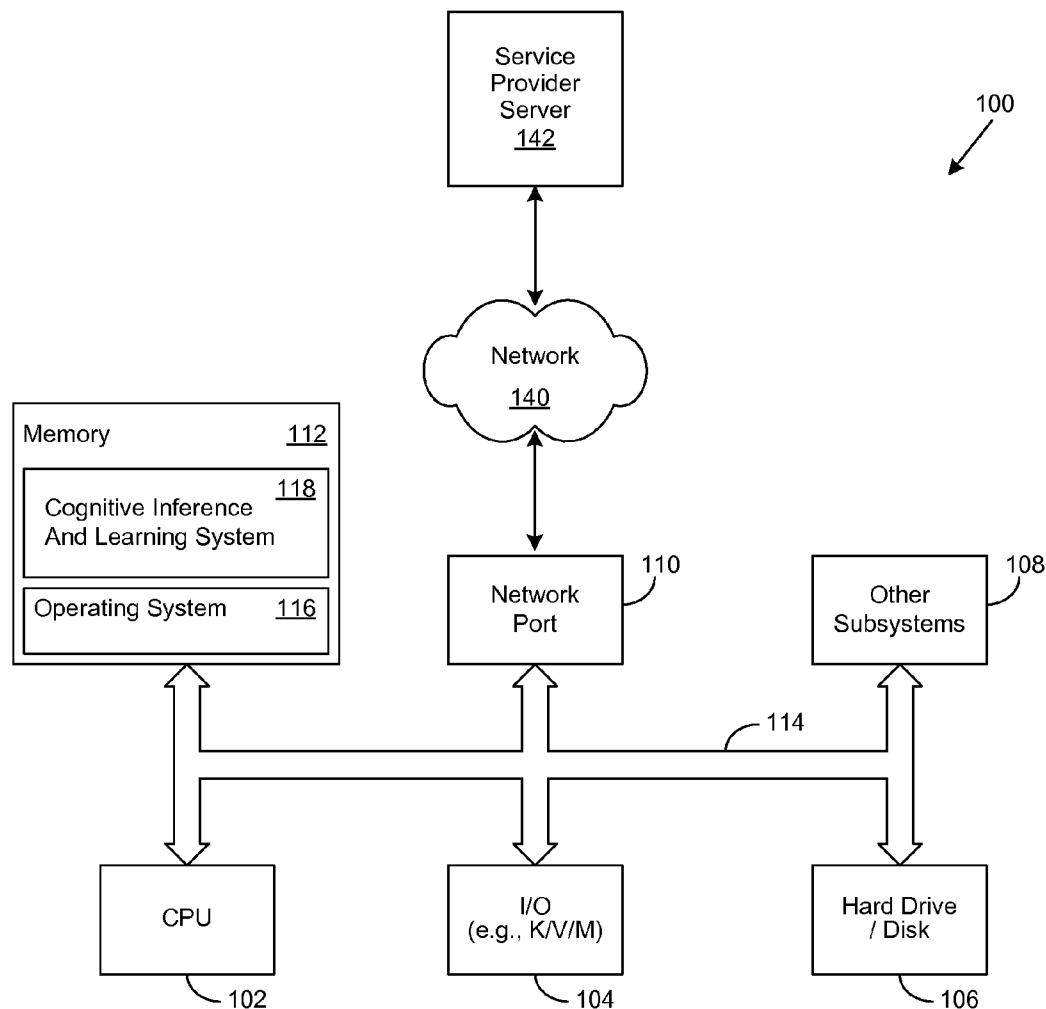




US 20150356414A1

(19) **United States**(12) **Patent Application Publication**
Sanchez et al.(10) **Pub. No.: US 2015/0356414 A1**(43) **Pub. Date: Dec. 10, 2015**(54) **UNIVERSAL KNOWLEDGE REPOSITORY**(71) Applicant: **Cognitive Scale, Inc.**, Austin, TX (US)(72) Inventors: **Matthew Sanchez**, Austin, TX (US);
Dilum Ranatunga, Austin, TX (US)(73) Assignee: **COGNITIVE SCALE, INC.**, Austin,
TX (US)(21) Appl. No.: **14/729,554**(22) Filed: **Jun. 3, 2015****Related U.S. Application Data**(60) Provisional application No. 62/009,626, filed on Jun.
9, 2014.**Publication Classification**(51) **Int. Cl.**
G06N 5/02 (2006.01)
G06F 17/30 (2006.01)(52) **U.S. Cl.**CPC **G06N 5/022** (2013.01); **G06F 17/30893**
(2013.01)(57) **ABSTRACT**

A system comprising: a processor; a data bus coupled to the processor; and a non-transitory, computer-readable storage medium embodying computer program code, the non-transitory, computer-readable storage medium being coupled to the data bus. The computer program code interacting with a plurality of computer operations and comprising instructions executable by the processor and configured for: receiving streams of data from a plurality of data sources; processing the streams of data from the plurality of data sources, the processing the streams of data from the plurality of data sources performing data enriching to provide data enriched data streams; and, storing the data streams and the data enriched data streams within the universal knowledge repository as a collection of knowledge elements.



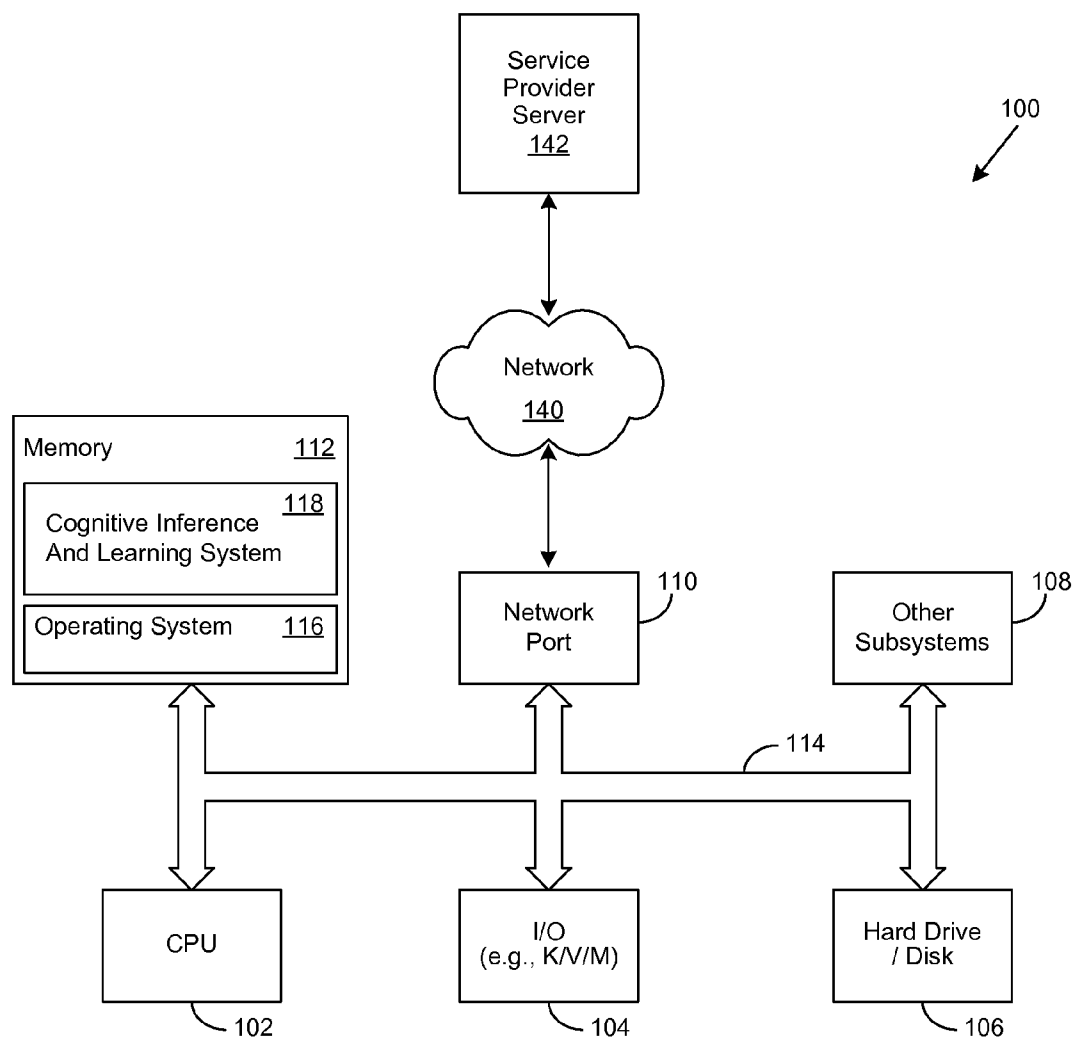


Figure 1

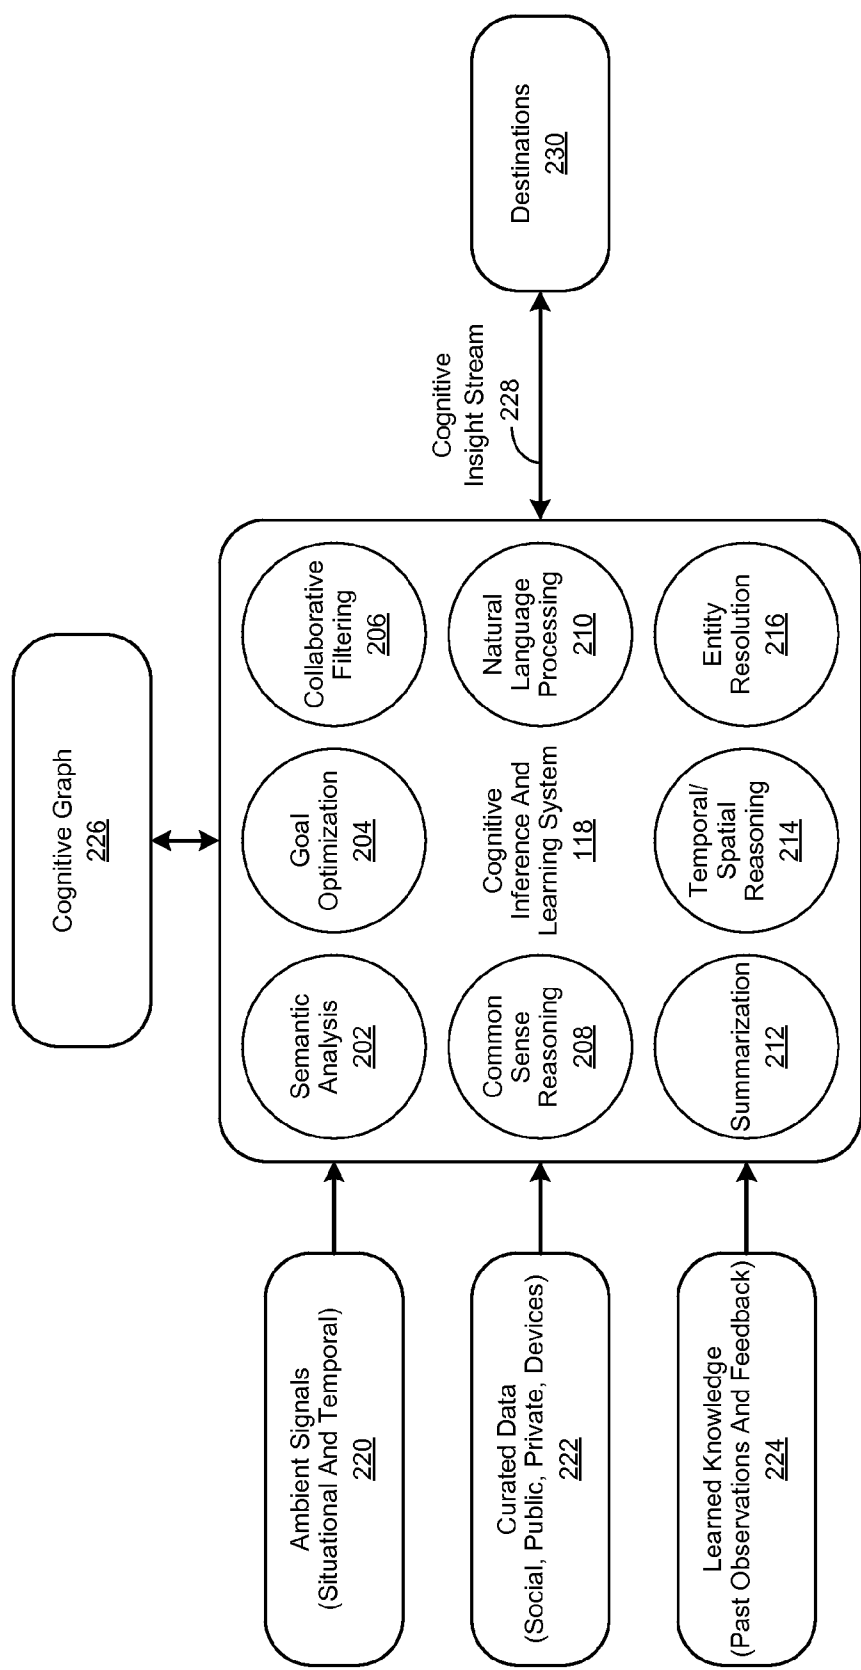


Figure 2

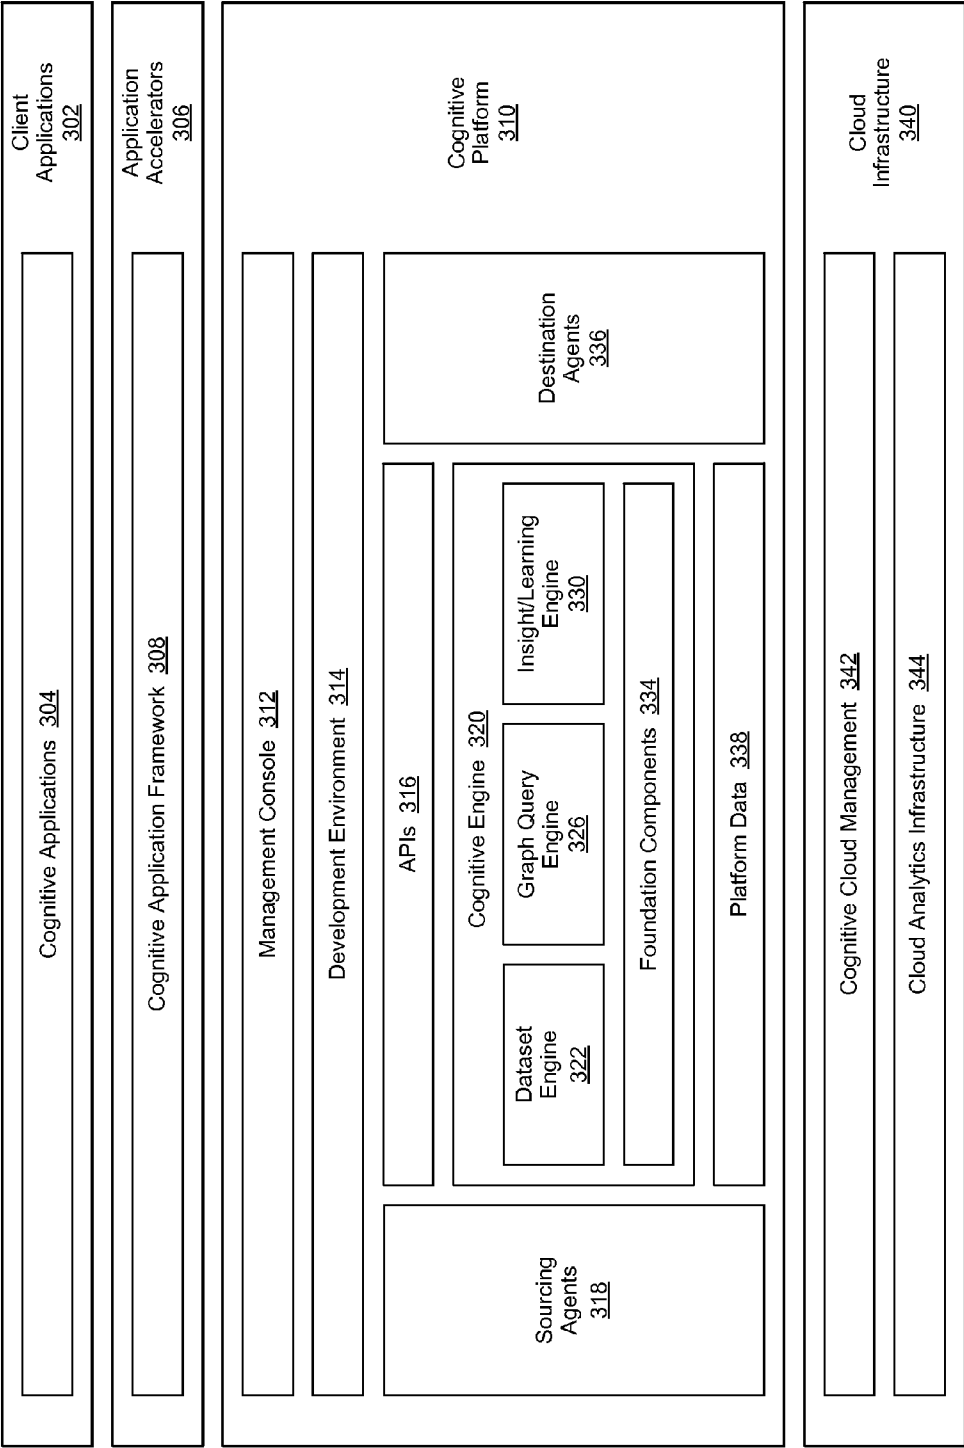
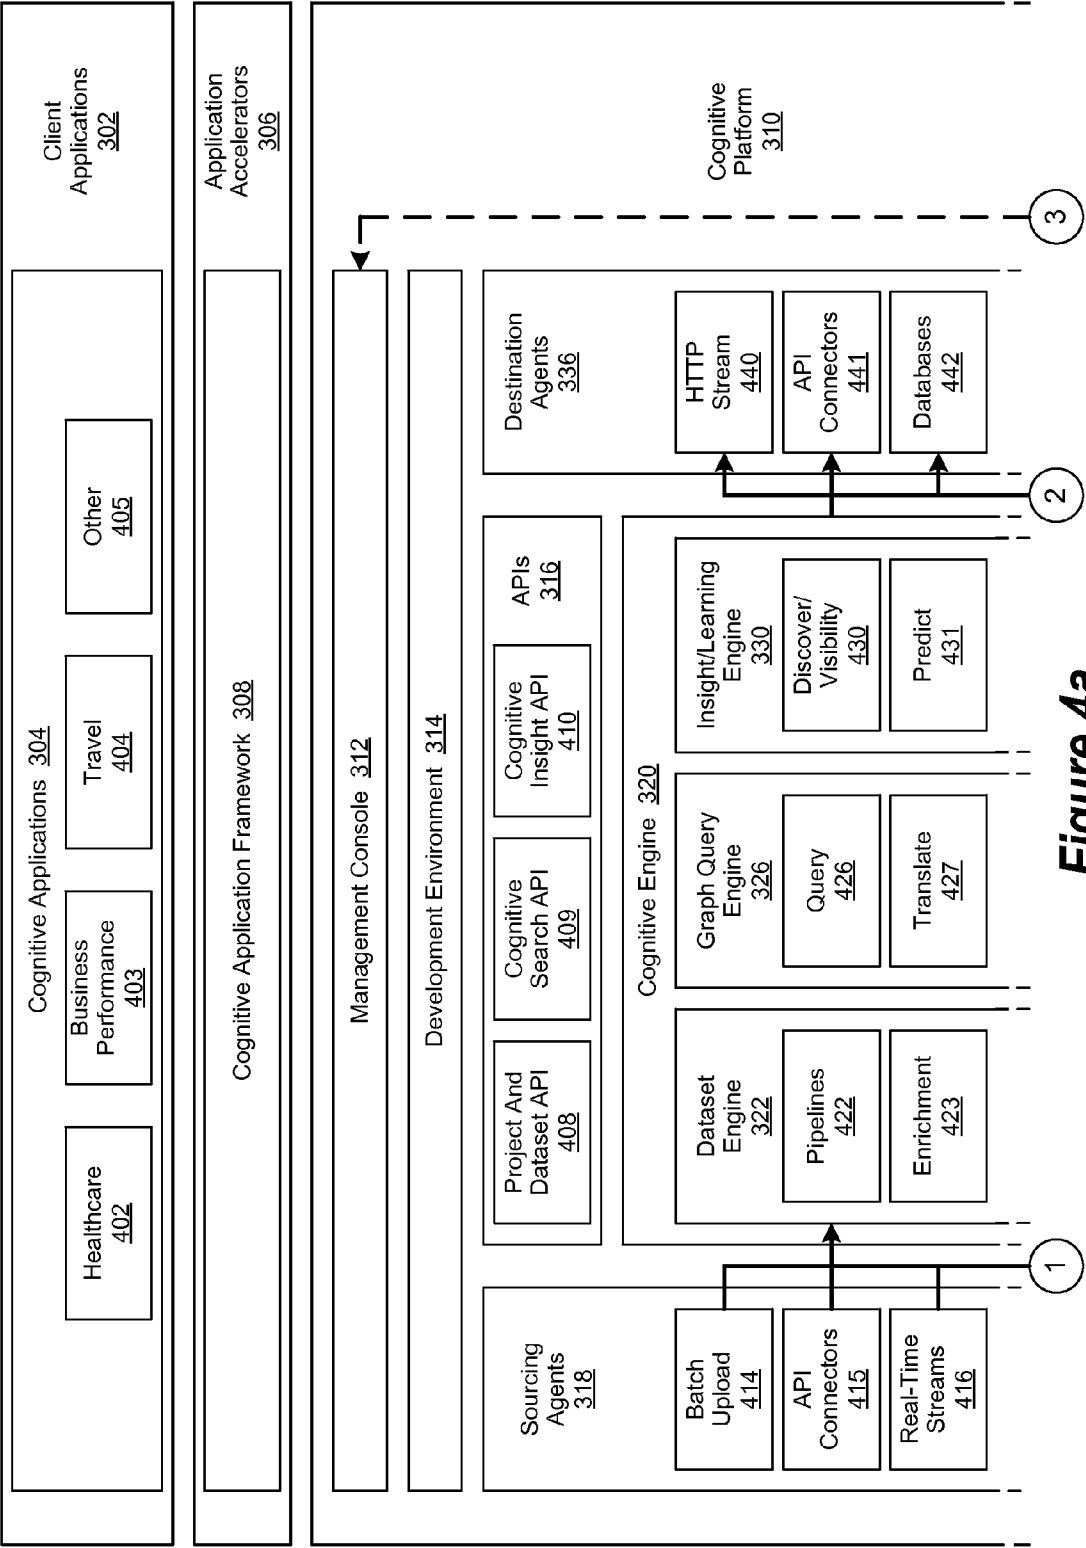


Figure 3



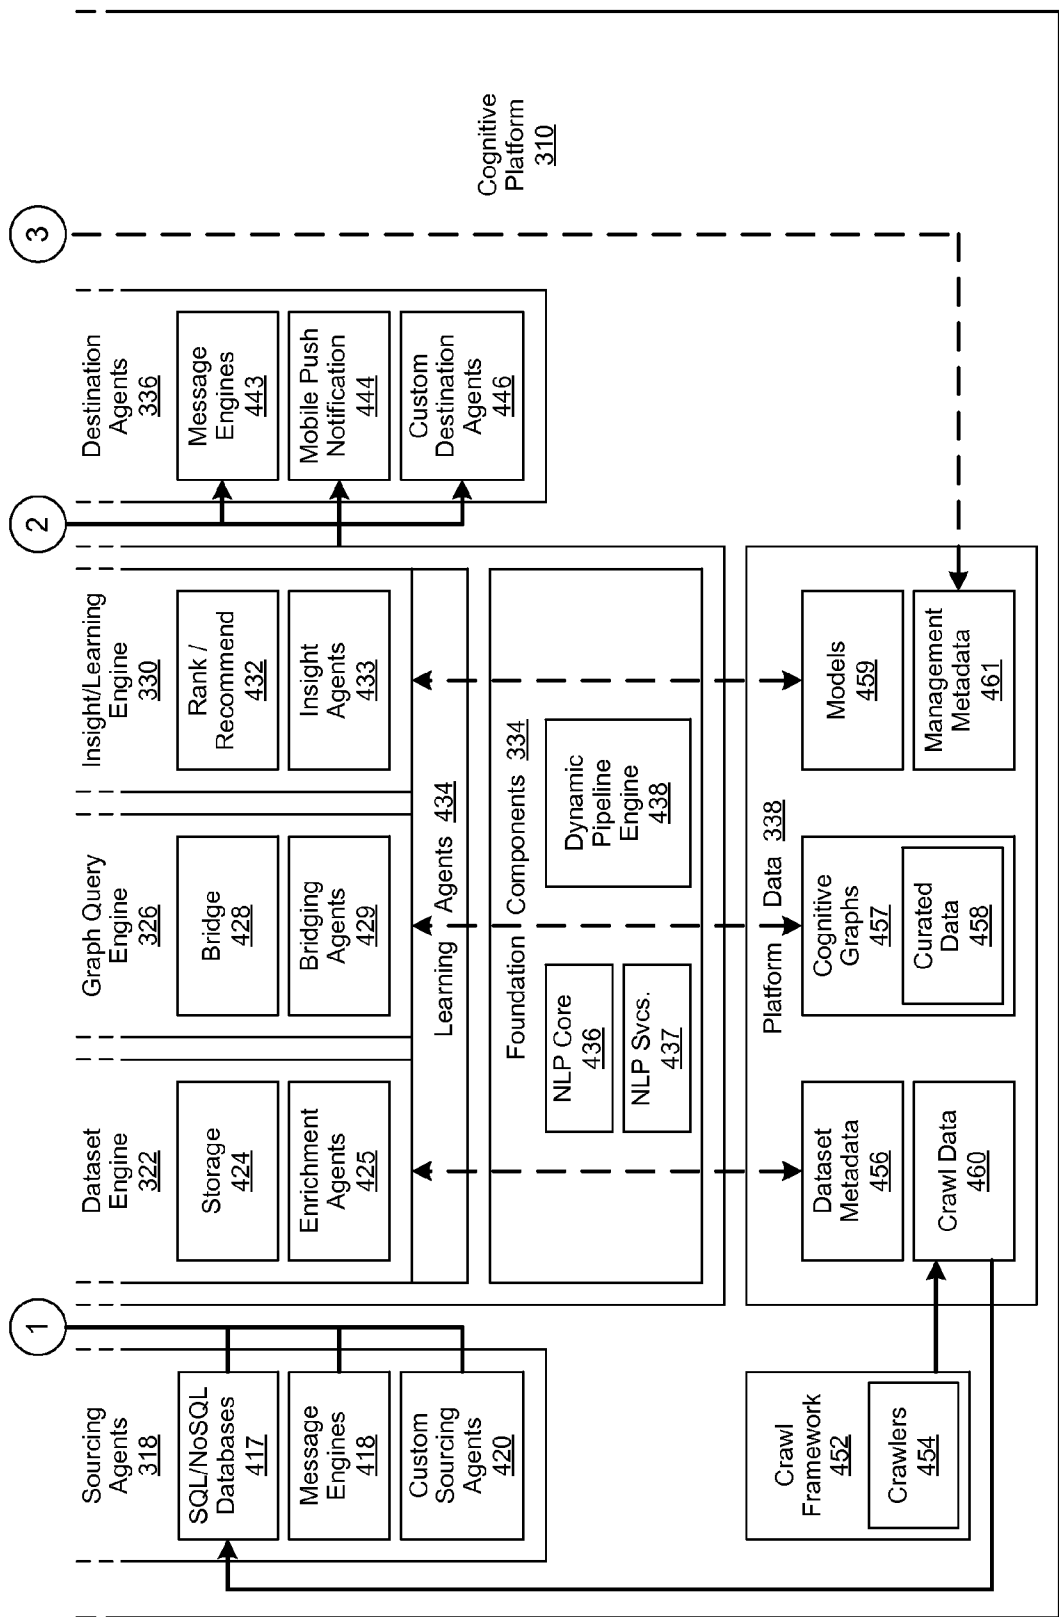


Figure 4b

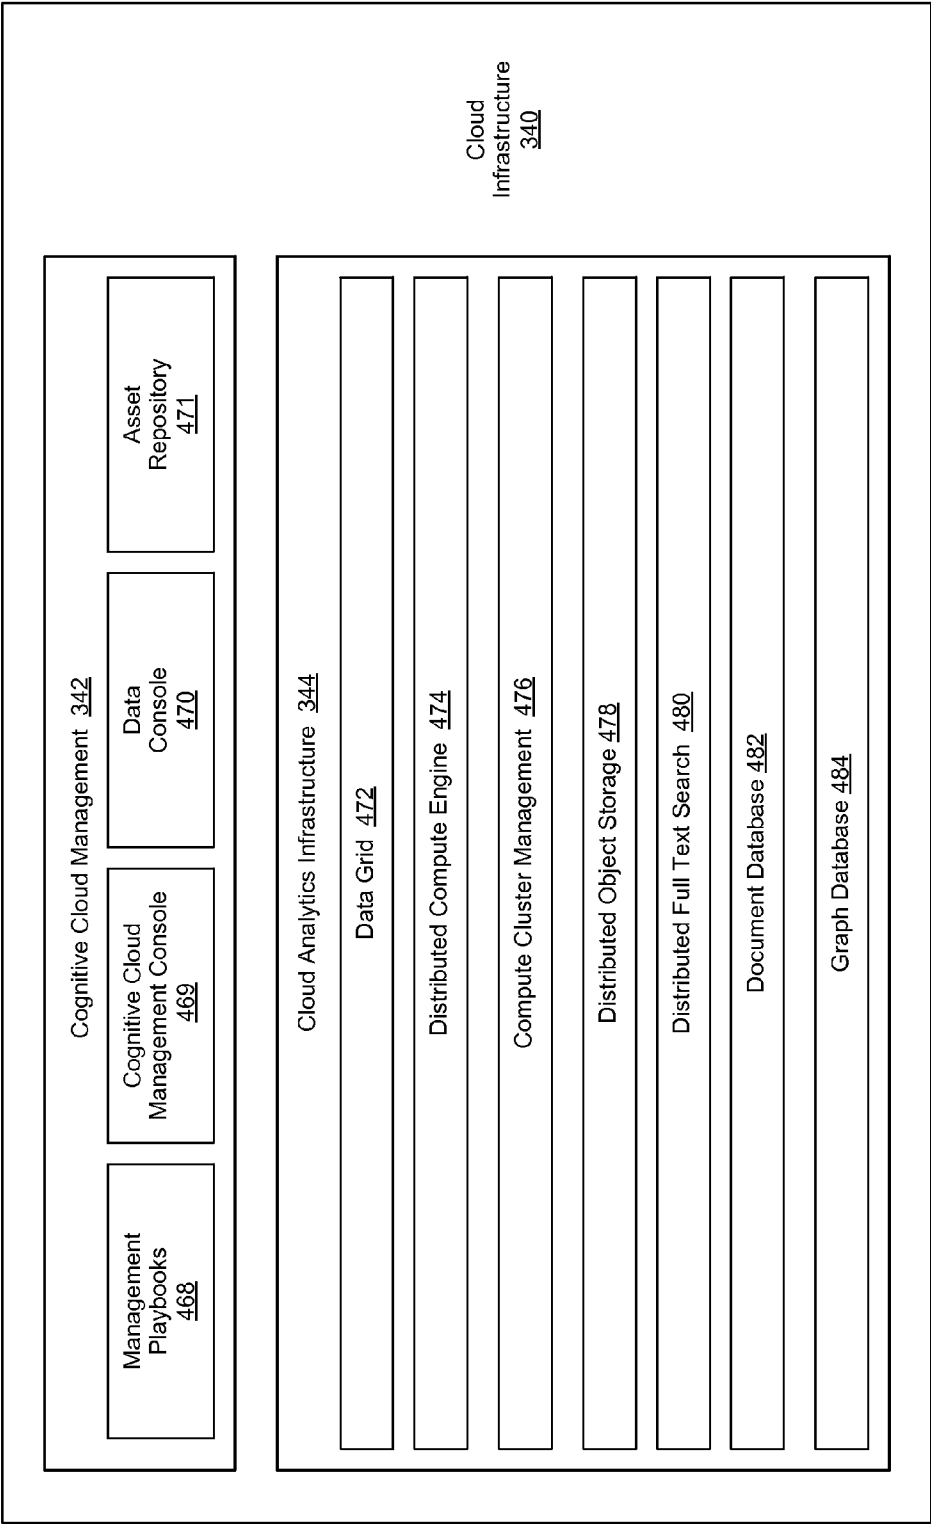


Figure 4c

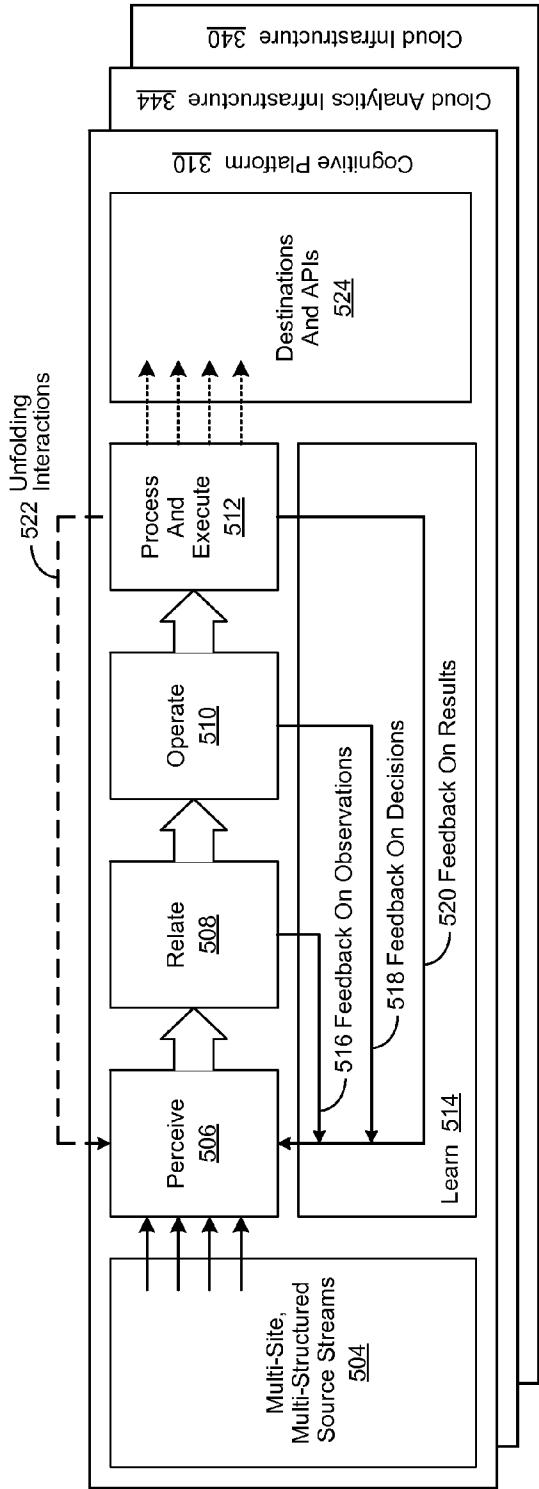


Figure 5

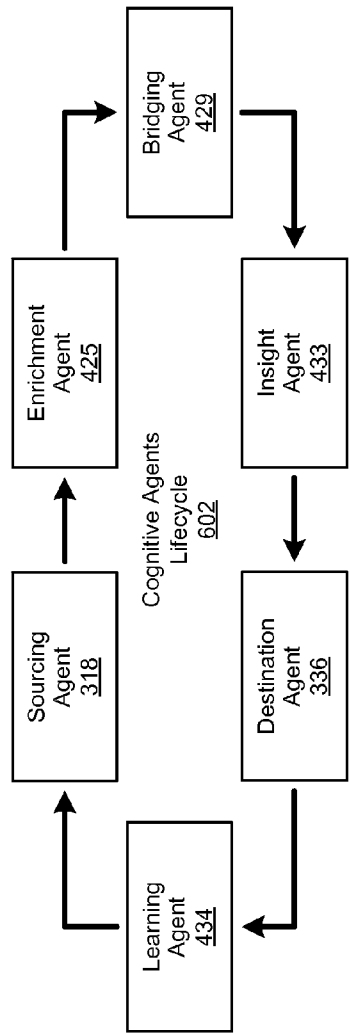


Figure 6

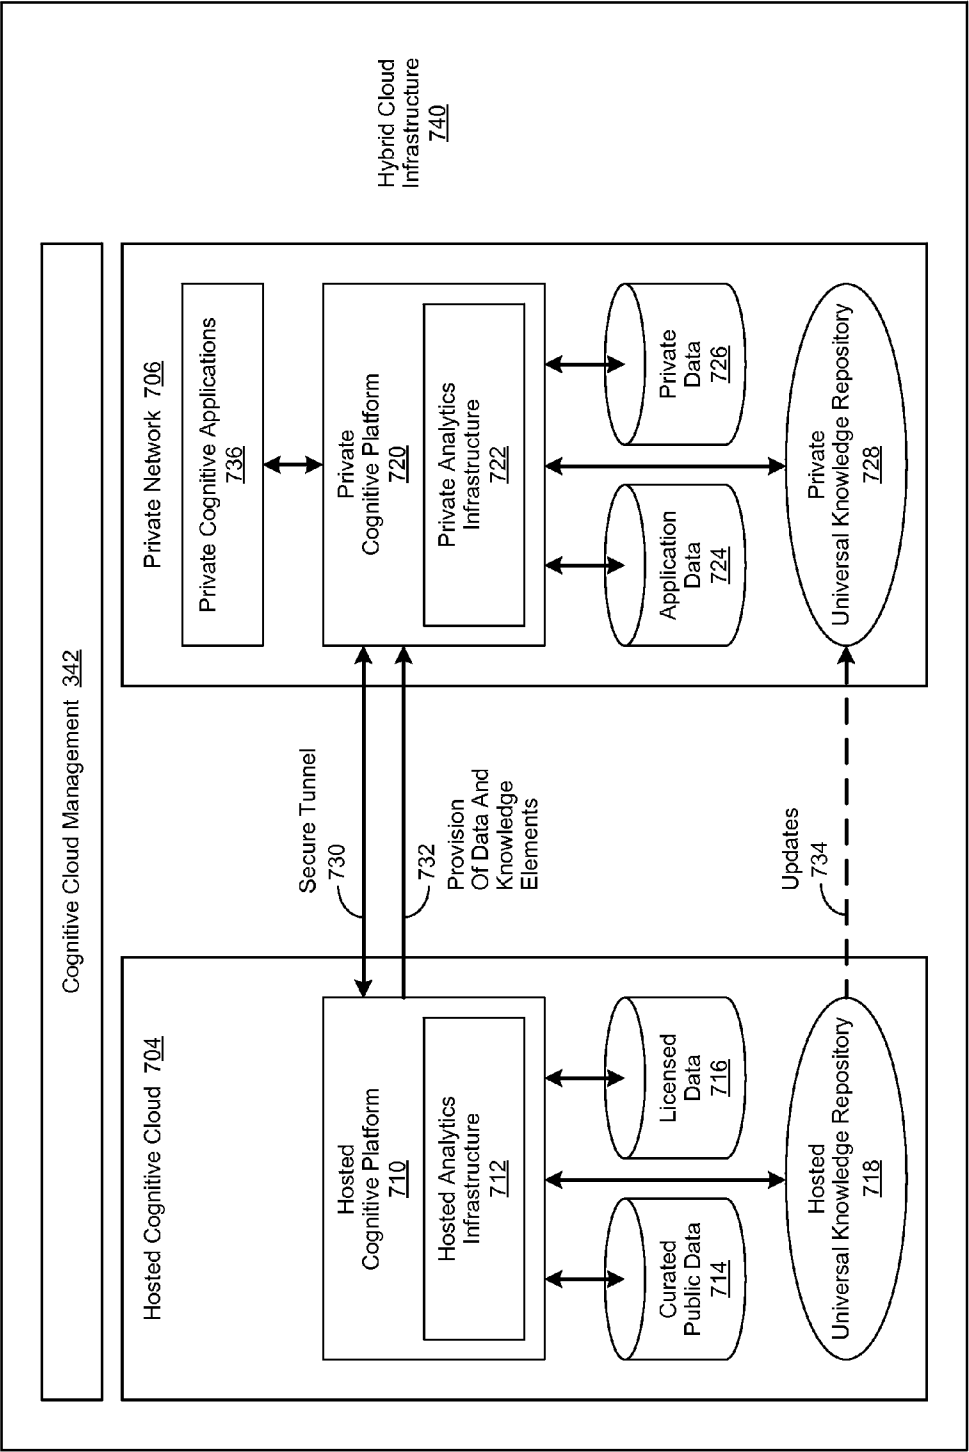


Figure 7

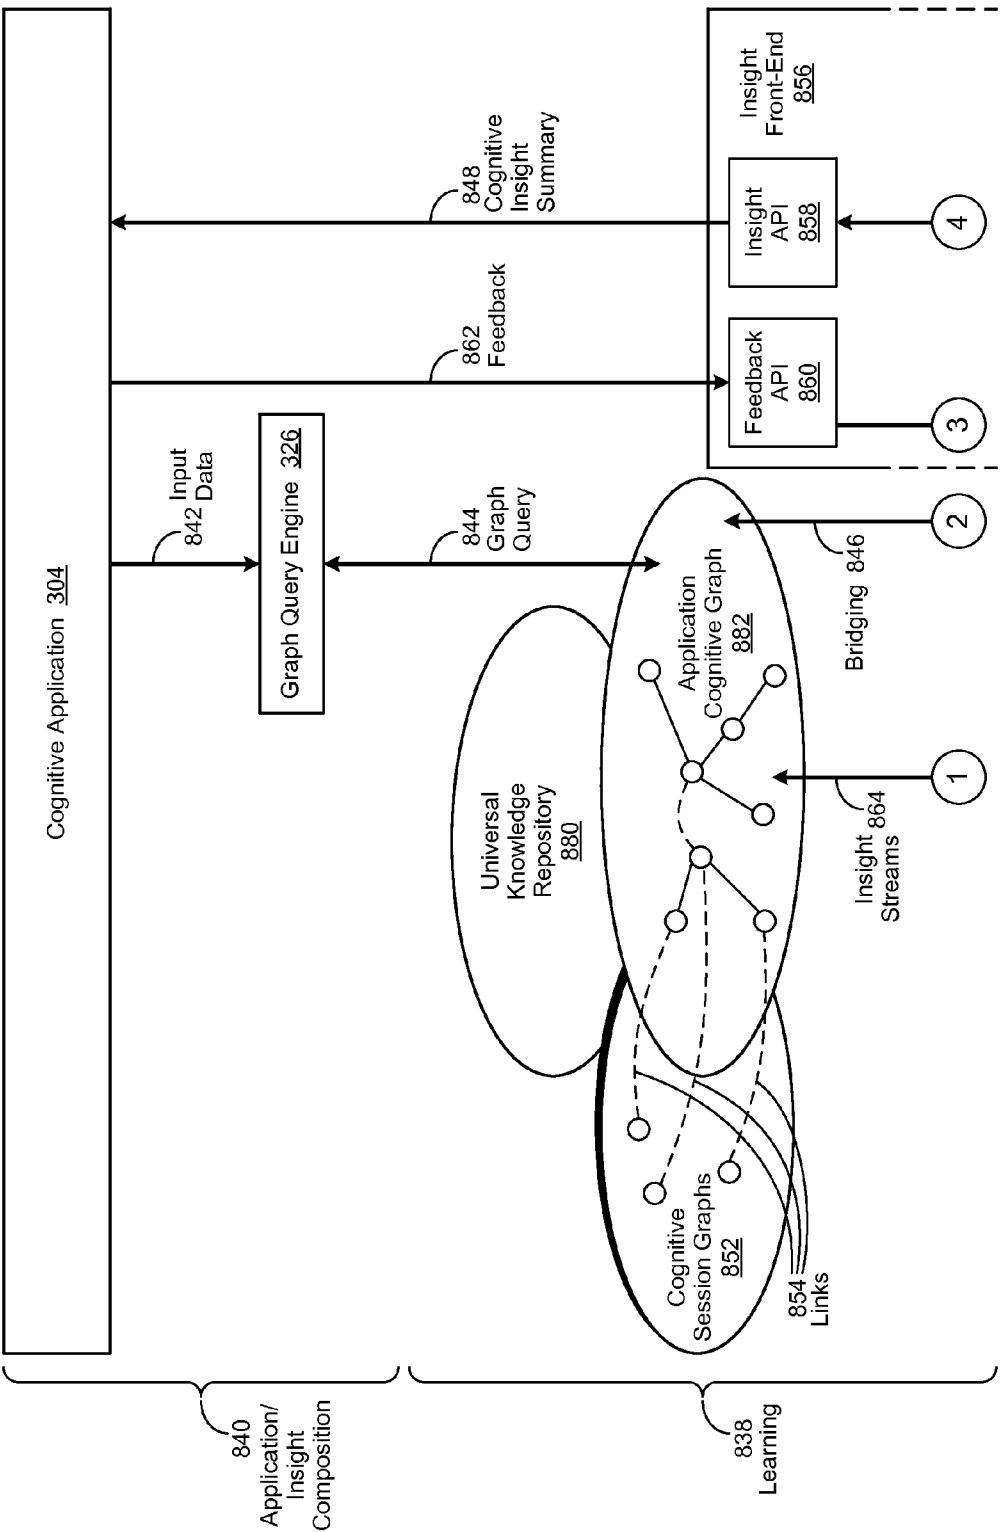


Figure 8a

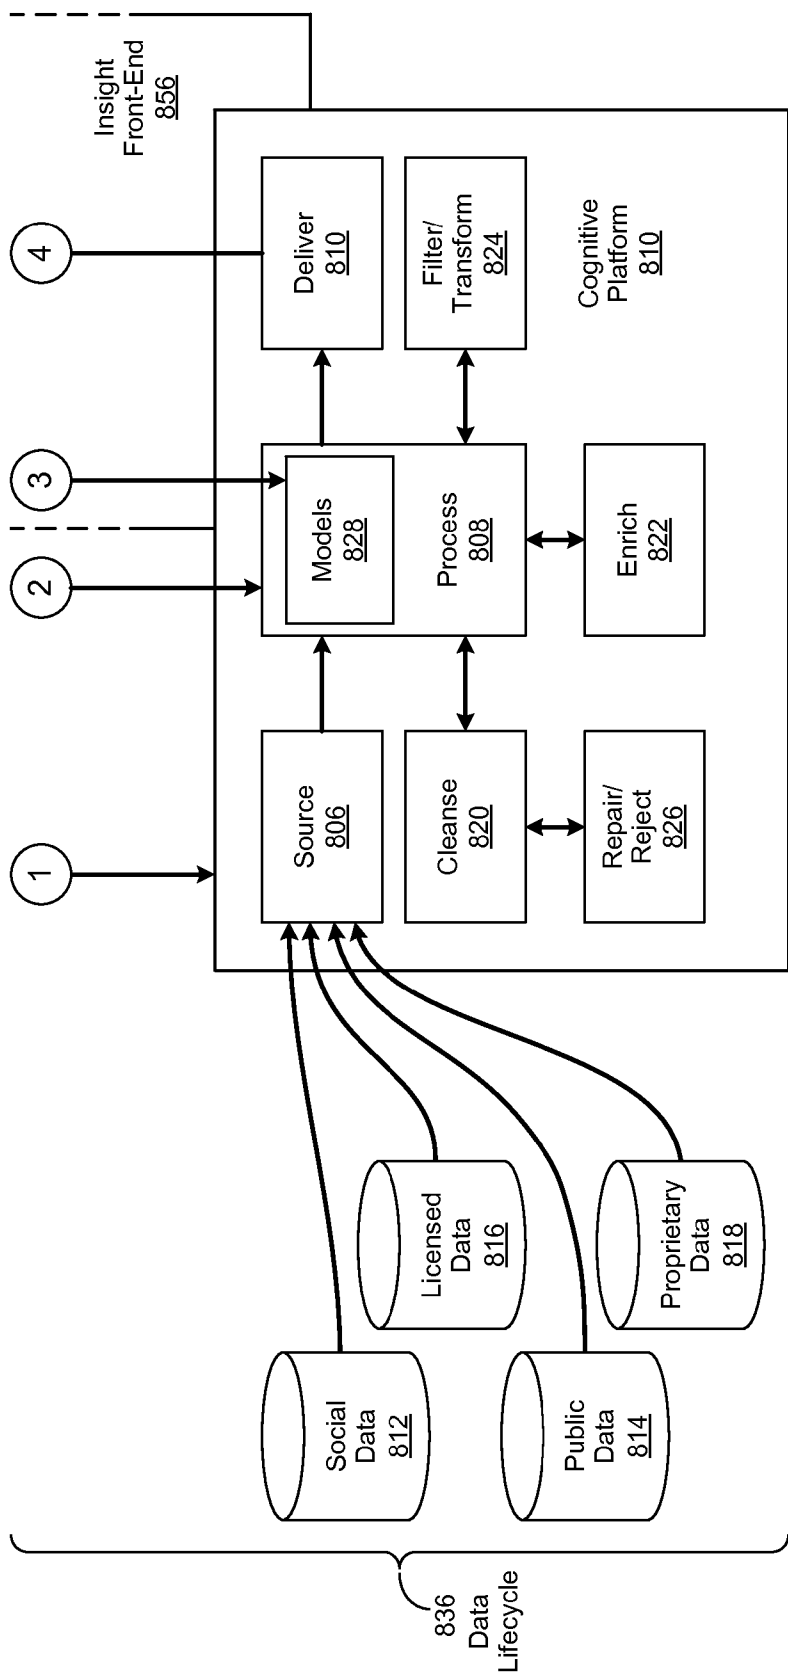


Figure 8b

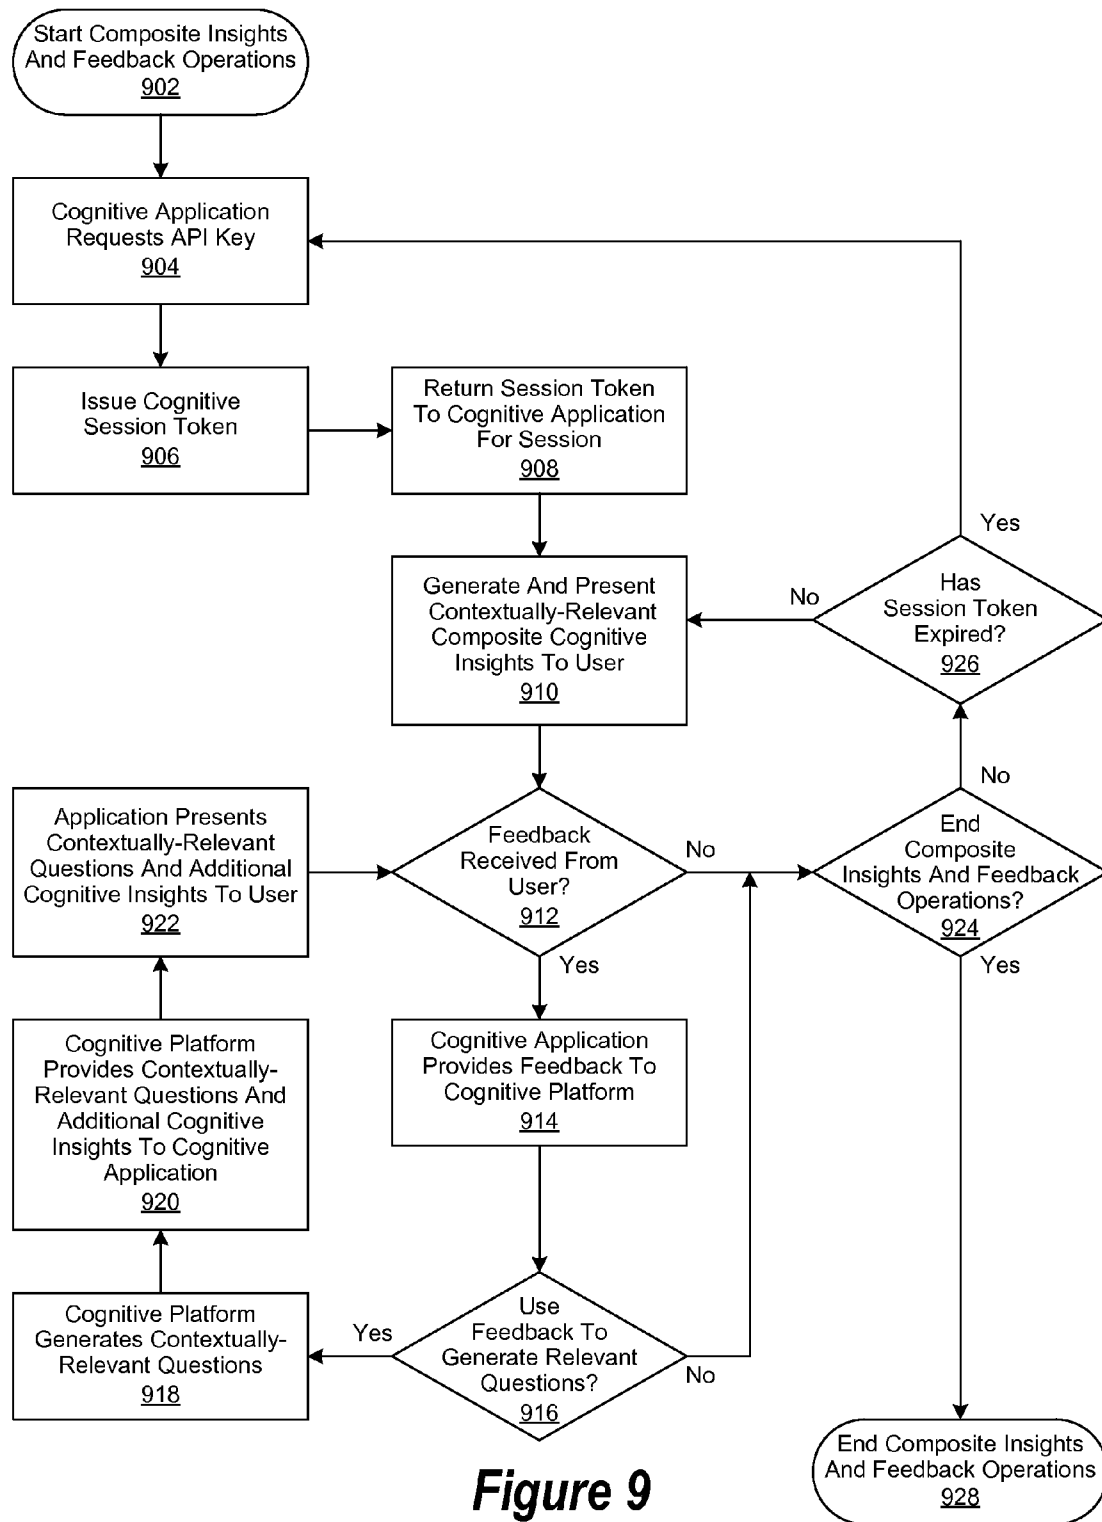


Figure 9

UNIVERSAL KNOWLEDGE REPOSITORY

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims the benefit under 35 U.S.C. §119(e) of U.S. Provisional Application No. 62/009,626, filed Jun. 9, 2014, entitled “Cognitive Information Processing System Environment.” U.S. Provisional Application No. 62/009,626 includes exemplary systems and methods and is incorporated by reference in its entirety.

BACKGROUND OF THE INVENTION

[0002] 1. Field of the Invention

[0003] The present invention relates in general to the field of computers and similar technologies, and in particular to software utilized in this field. Still more particularly, it relates to a method, system and computer-usable medium for generating and using a universal knowledge repository.

[0004] 2. Description of the Related Art

[0005] In general, “big data” refers to a collection of datasets so large and complex that they become difficult to process using typical database management tools and traditional data processing approaches. These datasets can originate from a wide variety of sources, including computer systems, mobile devices, credit card transactions, television broadcasts, and medical equipment, as well as infrastructures associated with cities, sensor-equipped buildings and factories, and transportation systems. Challenges commonly associated with big data, which may be a combination of structured, unstructured, and semi-structured data, include its capture, curation, storage, search, sharing, analysis and visualization. In combination, these challenges make it difficult to efficiently process large quantities of data within tolerable time intervals.

[0006] Nonetheless, big data analytics hold the promise of extracting insights by uncovering difficult-to-discover patterns and connections, as well as providing assistance in making complex decisions by analyzing different and potentially conflicting options. As such, individuals and organizations alike can be provided new opportunities to innovate, compete, and capture value.

[0007] One aspect of big data is “dark data,” which generally refers to data that is either not collected, neglected, or underutilized. Examples of data that is not currently being collected includes location data prior to the emergence of companies such as Foursquare or social data prior to the advent companies such as Facebook. An example of data that is being collected, but is difficult to access at the right time and place, includes data associated with the side effects of certain spider bites while on a camping trip. As another example, data that is collected and available, but has not yet been productized or fully utilized, may include disease insights from population-wide healthcare records and social media feeds. As a result, a case can be made that dark data may in fact be of higher value than big data in general, especially as it can likely provide actionable insights when it is combined with readily-available data.

SUMMARY OF THE INVENTION

[0008] A method, system and computer-usable medium are disclosed for cognitive inference and learning operations.

[0009] In one embodiment, the invention relates to a system comprising: a processor; a data bus coupled to the processor;

and a non-transitory, computer-readable storage medium embodying computer program code, the non-transitory, computer-readable storage medium being coupled to the data bus. The computer program code interacting with a plurality of computer operations and comprising instructions executable by the processor and configured for: receiving streams of data from a plurality of data sources; processing the streams of data from the plurality of data sources, the processing the streams of data from the plurality of data sources performing data enriching to provide data enriched data streams; and, storing the data streams and the data enriched data streams within the universal knowledge repository as a collection of knowledge elements.

[0010] In another embodiment, the invention relates to a non-transitory, computer-readable storage medium embodying computer program code, the computer program code comprising computer executable instructions configured for: receiving streams of data from a plurality of data sources; processing the streams of data from the plurality of data sources, the processing the streams of data from the plurality of data sources performing data enriching to provide data enriched data streams; and, storing the data streams and the data enriched data streams within the universal knowledge repository as a collection of knowledge elements.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The present invention may be better understood, and its numerous objects, features and advantages made apparent to those skilled in the art by referencing the accompanying drawings. The use of the same reference number throughout the several figures designates a like or similar element.

[0012] FIG. 1 depicts an exemplary client computer in which the present invention may be implemented;

[0013] FIG. 2 is a simplified block diagram of a cognitive inference and learning system (CILS);

[0014] FIG. 3 is a simplified block diagram of a CILS reference model implemented in accordance with an embodiment of the invention;

[0015] FIGS. 4a through 4c depict additional components of the CILS reference model shown in FIG. 3;

[0016] FIG. 5 is a simplified process diagram of CILS operations;

[0017] FIG. 6 depicts the lifecycle of CILS agents implemented to perform CILS operations;

[0018] FIG. 7 is a simplified block diagram of a plurality of cognitive platforms implemented in a hybrid cloud environment;

[0019] FIGS. 8a and 8b show a simplified process flow diagram of composite cognitive insight generation and feedback operations; and

[0020] FIG. 9 is a generalized flowchart of the performance of composite cognitive insight generation and feedback operations.

DETAILED DESCRIPTION

[0021] A method, system and computer-usable medium are disclosed for cognitive inference and learning operations. The present invention may be a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

[0022] The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: 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), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[0023] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

[0024] Computer readable program instructions for carrying out operations of the present invention 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 Smalltalk, C++ or the like, and conventional procedural programming languages, such as the “C” programming language or similar programming languages. The computer readable program instructions 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 any type of network, including 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 embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program

instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

[0025] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[0026] These computer readable 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 instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0027] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0028] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block 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. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

[0029] FIG. 1 is a generalized illustration of an information processing system 100 that can be used to implement the system and method of the present invention. The information processing system 100 includes a processor (e.g., central processor unit or “CPU”) 102, input/output (I/O) devices 104, such as a display, a keyboard, a mouse, and associated controllers, a hard drive or disk storage 106, and various other subsystems 108. In various embodiments, the information processing system 100 also includes network port 110 oper-

able to connect to a network **140**, which is likewise accessible by a service provider server **142**. The information processing system **100** likewise includes system memory **112**, which is interconnected to the foregoing via one or more buses **114**. System memory **112** further comprises operating system (OS) **116** and in various embodiments may also comprise cognitive inference and learning system (CILS) **118**. In these and other embodiments, the CILS **118** may likewise comprise invention modules **120**. In one embodiment, the information processing system **100** is able to download the CILS **118** from the service provider server **142**. In another embodiment, the CILS **118** is provided as a service from the service provider server **142**.

[0030] In various embodiments, the CILS **118** is implemented to perform various cognitive computing operations described in greater detail herein. As used herein, cognitive computing broadly refers to a class of computing involving self-learning systems that use techniques such as spatial navigation, machine vision, and pattern recognition to increasingly mimic the way the human brain works. To be more specific, earlier approaches to computing typically solved problems by executing a set of instructions codified within software. In contrast, cognitive computing approaches are data-driven, sense-making, insight-extracting, problem-solving systems that have more in common with the structure of the human brain than with the architecture of contemporary, instruction-driven computers.

[0031] To further differentiate these distinctions, traditional computers must first be programmed by humans to perform specific tasks, while cognitive systems learn from their interactions with data and humans alike, and in a sense, program themselves to perform new tasks. To summarize the difference between the two, traditional computers are designed to calculate rapidly. Cognitive systems are designed to quickly draw inferences from data and gain new knowledge.

[0032] Cognitive systems achieve these abilities by combining various aspects of artificial intelligence, natural language processing, dynamic learning, and hypothesis generation to render vast quantities of intelligible data to assist humans in making better decisions. As such, cognitive systems can be characterized as having the ability to interact naturally with people to extend what either humans, or machines, could do on their own. Furthermore, they are typically able to process natural language, multi-structured data, and experience much in the same way as humans. Moreover, they are also typically able to learn a knowledge domain based upon the best available data and get better, and more immersive, over time.

[0033] It will be appreciated that more data is currently being produced every day than was recently produced by human beings from the beginning of recorded time. Deep within this ever-growing mass of data is a class of data known as “dark data,” which includes neglected information, ambient signals, and insights that can assist organizations and individuals in augmenting their intelligence and deliver actionable insights through the implementation of cognitive applications. As used herein, cognitive applications, or “cognitive apps,” broadly refer to cloud-based, big data interpretive applications that learn from user engagement and data interactions. Such cognitive applications extract patterns and insights from dark data sources that are currently almost completely opaque. Examples of such dark data include disease insights from population-wide healthcare records and

social media feeds, or from new sources of information, such as sensors monitoring pollution in delicate marine environments.

[0034] Over time, it is anticipated that cognitive applications will fundamentally change the ways in which many organizations operate as they invert current issues associated with data volume and variety to enable a smart, interactive data supply chain. Ultimately, cognitive applications hold the promise of receiving a user query and immediately providing a data-driven answer from a masked data supply chain in response. As they evolve, it is likewise anticipated that cognitive applications may enable a new class of “sixth sense” applications that intelligently detect and learn from relevant data and events to offer insights, predictions and advice rather than wait for commands. Just as web and mobile applications changed the way people access data, cognitive applications may change the way people listen to, and become empowered by, multi-structured data such as emails, social media feeds, doctors notes, transaction records, and call logs.

[0035] However, the evolution of such cognitive applications has associated challenges, such as how to detect events, ideas, images, and other content that may be of interest. For example, assuming that the role and preferences of a given user are known, how is the most relevant information discovered, prioritized, and summarized from large streams of multi-structured data such as news feeds, blogs, social media, structured data, and various knowledge bases? To further the example, what can a healthcare executive be told about their competitor’s market share? Other challenges include the creation of a contextually-appropriate visual summary of responses to questions or queries.

[0036] FIG. 2 is a simplified block diagram of a cognitive inference and learning system (CILS) implemented in accordance with an embodiment of the invention. In various embodiments, the CILS **118** is implemented to incorporate a variety of processes, including semantic analysis **202**, goal optimization **204**, collaborative filtering **206**, common sense reasoning **208**, natural language processing **210**, summarization **212**, temporal/spatial reasoning **214**, and entity resolution **216** to generate cognitive insights.

[0037] As used herein, semantic analysis **202** broadly refers to performing various analysis operations to achieve a semantic level of understanding about language by relating syntactic structures. In various embodiments, various syntactic structures are related from the levels of phrases, clauses, sentences and paragraphs, to the level of the body of content as a whole and to its language-independent meaning. In certain embodiments, the semantic analysis **202** process includes processing a target sentence to parse it into its individual parts of speech, tag sentence elements that are related to predetermined items of interest, identify dependencies between individual words, and perform co-reference resolution. For example, if a sentence states that the author really likes the hamburgers served by a particular restaurant, then the name of the “particular restaurant” is co-referenced to “hamburgers.”

[0038] As likewise used herein, goal optimization **204** broadly refers to performing multi-criteria decision making operations to achieve a given goal or target objective. In various embodiments, one or more goal optimization **204** processes are implemented by the CILS **118** to define predetermined goals, which in turn contribute to the generation of a cognitive insight. For example, goals for planning a vacation trip may include low cost (e.g., transportation and

accommodations), location (e.g., by the beach), and speed (e.g., short travel time). In this example, it will be appreciated that certain goals may be in conflict with another. As a result, a cognitive insight provided by the CILS 118 to a traveler may indicate that hotel accommodations by a beach may cost more than they care to spend.

[0039] Collaborative filtering 206, as used herein, broadly refers to the process of filtering for information or patterns through the collaborative involvement of multiple agents, viewpoints, data sources, and so forth. The application of such collaborative filtering 206 processes typically involves very large and different kinds of data sets, including sensing and monitoring data, financial data, and user data of various kinds. Collaborative filtering 206 may also refer to the process of making automatic predictions associated with predetermined interests of a user by collecting preferences or other information from many users. For example, if person 'A' has the same opinion as a person 'B' for a given issue 'x', then an assertion can be made that person 'A' is more likely to have the same opinion as person 'B' opinion on a different issue 'y' than to have the same opinion on issue 'y' as a randomly chosen person. In various embodiments, the collaborative filtering 206 process is implemented with various recommendation engines familiar to those of skill in the art to make recommendations.

[0040] As used herein, common sense reasoning 208 broadly refers to simulating the human ability to make deductions from common facts they inherently know. Such deductions may be made from inherent knowledge about the physical properties, purpose, intentions and possible behavior of ordinary things, such as people, animals, objects, devices, and so on. In various embodiments, common sense reasoning 208 processes are implemented to assist the CILS 118 in understanding and disambiguating words within a predetermined context. In certain embodiments, the common sense reasoning 208 processes are implemented to allow the CILS 118 to generate text or phrases related to a target word or phrase to perform deeper searches for the same terms. It will be appreciated that if the context of a word is better understood, then a common sense understanding of the word can then be used to assist in finding better or more accurate information. In certain embodiments, this better or more accurate understanding of the context of a word, and its related information, allows the CILS 118 to make more accurate deductions, which are in turn used to generate cognitive insights.

[0041] As likewise used herein, natural language processing (NLP) 210 broadly refers to interactions with a system, such as the CILS 118, through the use of human, or natural, languages. In various embodiments, various NLP 210 processes are implemented by the CILS 118 to achieve natural language understanding, which enables it to not only derive meaning from human or natural language input, but to also generate natural language output.

[0042] Summarization 212, as used herein, broadly refers to processing a set of information, organizing and ranking it, and then generating a corresponding summary. As an example, a news article may be processed to identify its primary topic and associated observations, which are then extracted, ranked, and then presented to the user. As another example, page ranking operations may be performed on the same news article to identify individual sentences, rank them, order them, and determine which of the sentences are most impactful in describing the article and its content. As yet another example, a structured data record, such as a patient's

electronic medical record (EMR), may be processed using the summarization 212 process to generate sentences and phrases that describes the content of the EMR. In various embodiments, various summarization 212 processes are implemented by the CILS 118 to generate summarizations of content streams, which are in turn used to generate cognitive insights.

[0043] As used herein, temporal/spatial reasoning 214 broadly refers to reasoning based upon qualitative abstractions of temporal and spatial aspects of common sense knowledge, described in greater detail herein. For example, it is not uncommon for a predetermined set of data to change over time. Likewise, other attributes, such as its associated meta-data, may likewise change over time. As a result, these changes may affect the context of the data. To further the example, the context of asking someone what they believe they should be doing at 3:00 in the afternoon during the workday while they are at work may be quite different than asking the same user the same question at 3:00 on a Sunday afternoon when they are at home. In various embodiments, various temporal/spatial reasoning 214 processes are implemented by the CILS 118 to determine the context of queries, and associated data, which are in turn used to generate cognitive insights.

[0044] As likewise used herein, entity resolution 216 broadly refers to the process of finding elements in a set of data that refer to the same entity across different data sources (e.g., structured, non-structured, streams, devices, etc.), where the target entity does not share a common identifier. In various embodiments, the entity resolution 216 process is implemented by the CILS 118 to identify significant nouns, adjectives, phrases or sentence elements that represent various predetermined entities within one or more domains. From the foregoing, it will be appreciated that the implementation of one or more of the semantic analysis 202, goal optimization 204, collaborative filtering 206, common sense reasoning 208, natural language processing 210, summarization 212, temporal/spatial reasoning 214, and entity resolution 216 processes by the CILS 118 can facilitate the generation of a semantic, cognitive model.

[0045] In various embodiments, the CILS 118 receives ambient signals 220, curated data 222, and learned knowledge, which is then processed by the CILS 118 to generate one or more cognitive graphs 226. In turn, the one or more cognitive graphs 226 are further used by the CILS 118 to generate cognitive insight streams, which are then delivered to one or more destinations 230, as described in greater detail herein.

[0046] As used herein, ambient signals 220 broadly refer to input signals, or other data streams, that may contain data providing additional insight or context to the curated data 222 and learned knowledge 224 received by the CILS 118. For example, ambient signals may allow the CILS 118 to understand that a user is currently using their mobile device, at location 'x', at time 'y', doing activity 'z'. To further the example, there is a difference between the user using their mobile device while they are on an airplane versus using their mobile device after landing at an airport and walking between one terminal and another. To extend the example even further, ambient signals may add additional context, such as the user is in the middle of a three leg trip and has two hours before their next flight. Further, they may be in terminal A1, but their next flight is out of C1, it is lunchtime, and they want to know the best place to eat. Given the available time the user has,

their current location, restaurants that are proximate to their predicted route, and other factors such as food preferences, the CILS 118 can perform various cognitive operations and provide a recommendation for where the user can eat.

[0047] In various embodiments, the curated data 222 may include structured, unstructured, social, public, private, streaming, device or other types of data described in greater detail herein. In certain embodiments, the learned knowledge 224 is based upon past observations and feedback from the presentation of prior cognitive insight streams and recommendations. In various embodiments, the learned knowledge 224 is provided via a feedback loop that provides the learned knowledge 224 in the form of a learning stream of data.

[0048] As likewise used herein, a cognitive graph 226 refers to a representation of expert knowledge, associated with individuals and groups over a period of time, to depict relationships between people, places, and things using words, ideas, audio and images. As such, it is a machine-readable formalism for knowledge representation that provides a common framework allowing data and knowledge to be shared and reused across user, application, organization, and community boundaries.

[0049] In various embodiments, the information contained in, and referenced by, a cognitive graph 226 is derived from many sources (e.g., public, private, social, device), such as curated data 222. In certain of these embodiments, the cognitive graph 226 assists in the identification and organization of information associated with how people, places and things are related to one other. In various embodiments, the cognitive graph 226 enables automated agents, described in greater detail herein, to access the Web more intelligently, enumerate inferences through utilization of curated, structured data 222, and provide answers to questions by serving as a computational knowledge engine.

[0050] In certain embodiments, the cognitive graph 226 not only elicits and maps expert knowledge by deriving associations from data, it also renders higher level insights and accounts for knowledge creation through collaborative knowledge modeling. In various embodiments, the cognitive graph 226 is a machine-readable, declarative memory system that stores and learns both episodic memory (e.g., specific personal experiences associated with an individual or entity), and semantic memory, which stores factual information (e.g., geo location of an airport or restaurant).

[0051] For example, the cognitive graph 226 may know that a given airport is a place, and that there is a list of related places such as hotels, restaurants and departure gates. Furthermore, the cognitive graph 226 may know that people such as business travelers, families and college students use the airport to board flights from various carriers, eat at various restaurants, or shop at certain retail stores. The cognitive graph 226 may also have knowledge about the key attributes from various retail rating sites that travelers have used to describe the food and their experience at various venues in the airport over the past six months.

[0052] In certain embodiments, the cognitive insight stream 228 is bidirectional, and supports flows of information both too and from destinations 230. In these embodiments, the first flow is generated in response to receiving a query, and subsequently delivered to one or more destinations 230. The second flow is generated in response to detecting information about a user of one or more of the destinations 230. Such use results in the provision of information to the CILS 118. In response, the CILS 118 processes that information, in the

context of what it knows about the user, and provides additional information to the user, such as a recommendation. In various embodiments, the cognitive insight stream 228 is configured to be provided in a “push” stream configuration familiar to those of skill in the art. In certain embodiments, the cognitive insight stream 228 is implemented to use natural language approaches familiar to skilled practitioners of the art to support interactions with a user.

[0053] In various embodiments, the cognitive insight stream 228 may include a stream of visualized insights. As used herein, visualized insights broadly refers to cognitive insights that are presented in a visual manner, such as a map, an infographic, images, and so forth. In certain embodiments, these visualized insights may include various cognitive insights, such as “What happened?”, “What do I know about it?”, “What is likely to happen next?”, or “What should I do about it?” In these embodiments, the cognitive insight stream is generated by various cognitive agents, which are applied to various sources, datasets, and cognitive graphs. As used herein, a cognitive agent broadly refers to a computer program that performs a task with minimum specific directions from users and learns from each interaction with data and human users.

[0054] In various embodiments, the CILS 118 delivers Cognition as a Service (CaaS). As such, it provides a cloud-based development and execution platform that allow various cognitive applications and services to function more intelligently and intuitively. In certain embodiments, cognitive applications powered by the CILS 118 are able to think and interact with users as intelligent virtual assistants. As a result, users are able to interact with such cognitive applications by asking them questions and giving them commands. In response, these cognitive applications will be able to assist the user in completing tasks and managing their work more efficiently.

[0055] In these and other embodiments, the CILS 118 can operate as an analytics platform to process big data, and dark data as well, to provide data analytics through a public, private or hybrid cloud environment. As used herein, cloud analytics broadly refers to a service model wherein data sources, data models, processing applications, computing power, analytic models, and sharing or storage of results are implemented within a cloud environment to perform one or more aspects of analytics.

[0056] In various embodiments, users submit queries and computation requests in a natural language format to the CILS 118. In response, they are provided with a ranked list of relevant answers and aggregated information with useful links and pertinent visualizations through a graphical representation. In these embodiments, the cognitive graph 226 generates semantic and temporal maps to reflect the organization of unstructured data and to facilitate meaningful learning from potentially millions of lines of text, much in the same way as arbitrary syllables strung together create meaning through the concept of language.

[0057] FIG. 3 is a simplified block diagram of a cognitive inference and learning system (CILS) reference model implemented in accordance with an embodiment of the invention. In this embodiment, the CILS reference model is associated with the CILS 118 shown in FIG. 2. As shown in FIG. 3, the CILS 118 includes client applications 302, application accelerators 306, a cognitive platform 310, and cloud infrastructure 340. In various embodiments, the client applications 302 include cognitive applications 304, which are implemented to

understand and adapt to the user, not the other way around, by natively accepting and understanding human forms of communication, such as natural language text, audio, images, video, and so forth.

[0058] In these and other embodiments, the cognitive applications 304 possess situational and temporal awareness based upon ambient signals from users and data, which facilitates understanding the user's intent, content, context and meaning to drive goal-driven dialogs and outcomes. Further, they are designed to gain knowledge over time from a wide variety of structured, non-structured, and device data sources, continuously interpreting and autonomously reprogramming themselves to better understand a given domain. As such, they are well-suited to support human decision making, by proactively providing trusted advice, offers and recommendations while respecting user privacy and permissions.

[0059] In various embodiments, the application accelerators 306 include a cognitive application framework 308. In certain embodiments, the application accelerators 306 and the cognitive application framework 308 support various plug-ins and components that facilitate the creation of client applications 302 and cognitive applications 304. In various embodiments, the application accelerators 306 include widgets, user interface (UI) components, reports, charts, and back-end integration components familiar to those of skill in the art.

[0060] As likewise shown in FIG. 3, the cognitive platform 310 includes a management console 312, a development environment 314, application program interfaces (APIs) 316, sourcing agents 318, a cognitive engine 320, destination agents 336, and platform data 338, all of which are described in greater detail herein. In various embodiments, the management console 312 is implemented to manage accounts and projects, along with user-specific metadata that is used to drive processes and operations within the cognitive platform 310 for a predetermined project.

[0061] In certain embodiments, the development environment 314 is implemented to create custom extensions to the CILS 118 shown in FIG. 2. In various embodiments, the development environment 314 is implemented for the development of a custom application, which may subsequently be deployed in a public, private or hybrid cloud environment. In certain embodiments, the development environment 314 is implemented for the development of a custom sourcing agent, a custom bridging agent, a custom destination agent, or various analytics applications or extensions.

[0062] In various embodiments, the APIs 316 are implemented to build and manage predetermined cognitive applications 304, described in greater detail herein, which are then executed on the cognitive platform 310 to generate cognitive insights. Likewise, the sourcing agents 318 are implemented in various embodiments to source a variety of multi-site, multi-structured source streams of data described in greater detail herein. In various embodiments, the cognitive engine 320 includes a dataset engine 322, a graph query engine 326, an insight/learning engine 330, and foundation components 334. In certain embodiments, the dataset engine 322 is implemented to establish and maintain a dynamic data ingestion and enrichment pipeline. In these and other embodiments, the dataset engine 322 may be implemented to orchestrate one or more sourcing agents 318 to source data. Once the data is sourced, the data set engine 322 performs data enriching and other data processing operations, described in greater detail

herein, and generates one or more sub-graphs that are subsequently incorporated into a target cognitive graph.

[0063] In various embodiments, the graph query engine 326 is implemented to receive and process queries such that they can be bridged into a cognitive graph, as described in greater detail herein, through the use of a bridging agent. In certain embodiments, the graph query engine 326 performs various natural language processing (NLP), familiar to skilled practitioners of the art, to process the queries. In various embodiments, the insight/learning engine 330 is implemented to encapsulate a predetermined algorithm, which is then applied to a cognitive graph to generate a result, such as a cognitive insight or a recommendation. In certain embodiments, one or more such algorithms may contribute to answering a specific question and provide additional cognitive insights or recommendations. In various embodiments, two or more of the dataset engine 322, the graph query engine 326, and the insight/learning engine 330 may be implemented to operate collaboratively to generate a cognitive insight or recommendation. In certain embodiments, one or more of the dataset engine 322, the graph query engine 326, and the insight/learning engine 330 may operate autonomously to generate a cognitive insight or recommendation.

[0064] The foundation components 334 shown in FIG. 3 include various reusable components, familiar to those of skill in the art, which are used in various embodiments to enable the dataset engine 322, the graph query engine 326, and the insight/learning engine 330 to perform their respective operations and processes. Examples of such foundation components 334 include natural language processing (NLP) components and core algorithms, such as cognitive algorithms.

[0065] In various embodiments, the platform data 338 includes various data repositories, described in greater detail herein, that are accessed by the cognitive platform 310 to generate cognitive insights. In various embodiments, the destination agents 336 are implemented to publish cognitive insights to a consumer of cognitive insight data. Examples of such consumers of cognitive insight data include target databases, business intelligence applications, and mobile applications. It will be appreciated that many such examples of cognitive insight data consumers are possible and the foregoing is not intended to limit the spirit, scope or intent of the invention. In various embodiments, as described in greater detail herein, the cloud infrastructure 340 includes cognitive cloud management 342 components and cloud analytics infrastructure components 344.

[0066] FIGS. 4a through 4c depict additional cognitive inference and learning system (CILS) components implemented in accordance with an embodiment of the CILS reference model shown in FIG. 3. In this embodiment, the CILS reference model includes client applications 302, application accelerators 306, a cognitive platform 310, and cloud infrastructure 340. As shown in FIG. 4a, the client applications 302 include cognitive applications 304. In various embodiments, the cognitive applications 304 are implemented natively accept and understand human forms of communication, such as natural language text, audio, images, video, and so forth. In certain embodiments, the cognitive applications 304 may include healthcare 402, business performance 403, travel 404, and various other 405 applications familiar to skilled practitioners of the art. As such, the foregoing is only

provided as examples of such cognitive applications 304 and is not intended to limit the intent, spirit or scope of the invention.

[0067] In various embodiments, the application accelerators 306 include a cognitive application framework 308. In certain embodiments, the application accelerators 308 and the cognitive application framework 308 support various plug-ins and components that facilitate the creation of client applications 302 and cognitive applications 304. In various embodiments, the application accelerators 306 include widgets, user interface (UI) components, reports, charts, and back-end integration components familiar to those of skill in the art. It will be appreciated that many such application accelerators 306 are possible and their provided functionality, selection, provision and support are a matter of design choice. As such, the application accelerators 306 described in greater detail herein are not intended to limit the spirit, scope or intent of the invention.

[0068] As shown in FIGS. 4a and 4b, the cognitive platform 310 includes a management console 312, a development environment 314, application program interfaces (APIs) 316, sourcing agents 318, a cognitive engine 320, destination agents 336, platform data 338, and a crawl framework 452. In various embodiments, the management console 312 is implemented to manage accounts and projects, along with management metadata 461 that is used to drive processes and operations within the cognitive platform 310 for a predetermined project.

[0069] In various embodiments, the management console 312 is implemented to run various services on the cognitive platform 310. In certain embodiments, the management console 312 is implemented to manage the configuration of the cognitive platform 310. In certain embodiments, the management console 312 is implemented to establish the development environment 314. In various embodiments, the management console 312 may be implemented to manage the development environment 314 once it is established. Skilled practitioners of the art will realize that many such embodiments are possible and the foregoing is not intended to limit the spirit, scope or intent of the invention.

[0070] In various embodiments, the development environment 314 is implemented to create custom extensions to the CILS 118 shown in FIG. 2. In these and other embodiments, the development environment 314 is implemented to support various programming languages, such as Python, Java, R, and others familiar to skilled practitioners of the art. In various embodiments, the development environment 314 is implemented to allow one or more of these various programming languages to create a variety of analytic models and applications. As an example, the development environment 314 may be implemented to support the R programming language, which in turn can be used to create an analytic model that is then hosted on the cognitive platform 310.

[0071] In certain embodiments, the development environment 314 is implemented for the development of various custom applications or extensions related to the cognitive platform 310, which may subsequently be deployed in a public, private or hybrid cloud environment. In various embodiments, the development environment 314 is implemented for the development of various custom sourcing agents 318, custom enrichment agents 425, custom bridging agents 429, custom insight agents 433, custom destination agents 336, and custom learning agents 434, which are described in greater detail herein.

[0072] In various embodiments, the APIs 316 are implemented to build and manage predetermined cognitive applications 304, described in greater detail herein, which are then executed on the cognitive platform 310 to generate cognitive insights. In these embodiments, the APIs 316 may include one or more of a project and dataset API 408, a cognitive search API 409, a cognitive insight API 410, and other APIs. The selection of the individual APIs 316 implemented in various embodiments is a matter design choice and the foregoing is not intended to limit the spirit, scope or intent of the invention.

[0073] In various embodiments, the project and dataset API 408 is implemented with the management console 312 to enable the management of a variety of data and metadata associated with various cognitive insight projects and user accounts hosted or supported by the cognitive platform 310. In one embodiment, the data and metadata managed by the project and dataset API 408 are associated with billing information familiar to those of skill in the art. In one embodiment, the project and dataset API 408 is used to access a data stream that is created, configured and orchestrated, as described in greater detail herein, by the dataset engine 322.

[0074] In various embodiments, the cognitive search API 409 uses natural language processes familiar to those of skill in the art to search a target cognitive graph. Likewise, the cognitive insight API 410 is implemented in various embodiments to configure the insight/learning engine 330 to provide access to predetermined outputs from one or more cognitive graph algorithms that are executing in the cognitive platform 310. In certain embodiments, the cognitive insight API 410 is implemented to subscribe to, or request, such predetermined outputs.

[0075] In various embodiments, the sourcing agents 318 may include a batch upload 414 agent, an API connectors 415 agent, a real-time streams 416 agent, a Structured Query Language (SQL)/Not Only SQL (NoSQL) databases 417 agent, a message engines 418 agent, and one or more custom sourcing 420 agents. Skilled practitioners of the art will realize that other types of sourcing agents 318 may be used in various embodiments and the foregoing is not intended to limit the spirit, scope or intent of the invention. In various embodiments, the sourcing agents 318 are implemented to source a variety of multi-site, multi-structured source streams of data described in greater detail herein. In certain embodiments, each of the sourcing agents 318 has a corresponding API.

[0076] In various embodiments, the batch uploading 414 agent is implemented for batch uploading of data to the cognitive platform 310. In these embodiments, the uploaded data may include a single data element, a single data record or file, or a plurality of data records or files. In certain embodiments, the data may be uploaded from more than one source and the uploaded data may be in a homogenous or heterogeneous form. In various embodiments, the API connectors 415 agent is implemented to manage interactions with one or more predetermined APIs that are external to the cognitive platform 310. As an example, Associated Press® may have their own API for news stories, Expedia® for travel information, or the National Weather Service for weather information. In these examples, the API connectors 415 agent would be implemented to determine how to respectively interact with each organization's API such that the cognitive platform 310 can receive information.

[0077] In various embodiments, the real-time streams 416 agent is implemented to receive various streams of data, such as social media streams (e.g., Twitter feeds) or other data streams (e.g., device data streams). In these embodiments, the streams of data are received in near-real-time. In certain embodiments, the data streams include temporal attributes. As an example, as data is added to a blog file, it is time-stamped to create temporal data. Other examples of a temporal data stream include Twitter feeds, stock ticker streams, device location streams from a device that is tracking location, medical devices tracking a patient's vital signs, and intelligent thermostats used to improve energy efficiency for homes.

[0078] In certain embodiments, the temporal attributes define a time window, which can be correlated to various elements of data contained in the stream. For example, as a given time window changes, associated data may have a corresponding change. In various embodiments, the temporal attributes do not define a time window. As an example, a social media feed may not have predetermined time windows, yet it is still temporal. As a result, the social media feed can be processed to determine what happened in the last 24 hours, what happened in the last hour, what happened in the last 15 minutes, and then determine related subject matter that is trending.

[0079] In various embodiments, the SQL/NoSQL databases 417 agent is implemented to interact with one or more target databases familiar to those of skill in the art. For example, the target database may include a SQL, NoSQL, delimited flat file, or other form of database. In various embodiments, the message engines 418 agent is implemented to provide data to the cognitive platform 310 from one or more message engines, such as a message queue (MQ) system, a message bus, a message broker, an enterprise service bus (ESB), and so forth. Skilled practitioners of the art will realize that there are many such examples of message engines with which the message engines 418 agent may interact and the foregoing is not intended to limit the spirit, scope or intent of the invention.

[0080] In various embodiments, the custom sourcing agents 420, which are purpose-built, are developed through the use of the development environment 314, described in greater detail herein. Examples of custom sourcing agents 420 include sourcing agents for various electronic medical record (EMR) systems at various healthcare facilities. Such EMR systems typically collect a variety of healthcare information, much of it the same, yet it may be collected, stored and provided in different ways. In this example, the custom sourcing agents 420 allow the cognitive platform 310 to receive information from each disparate healthcare source.

[0081] In various embodiments, the cognitive engine 320 includes a dataset engine 322, a graph engine 326, an insight/learning engine 330, learning agents 434, and foundation components 334. In these and other embodiments, the dataset engine 322 is implemented as described in greater detail to establish and maintain a dynamic data ingestion and enrichment pipeline. In various embodiments, the dataset engine 322 may include a pipelines 422 component, an enrichment 423 component, a storage component 424, and one or more enrichment agents 425.

[0082] In various embodiments, the pipelines 422 component is implemented to ingest various data provided by the sourcing agents 318. Once ingested, this data is converted by the pipelines 422 component into streams of data for process-

ing. In certain embodiments, these managed streams are provided to the enrichment 423 component, which performs data enrichment operations familiar to those of skill in the art. As an example, a data stream may be sourced from Associated Press® by a sourcing agent 318 and provided to the dataset engine 322. The pipelines 422 component receives the data stream and routes it to the enrichment 423 component, which then enriches the data stream by performing sentiment analysis, geotagging, and entity detection operations to generate an enriched data stream. In certain embodiments, the enrichment operations include filtering operations familiar to skilled practitioners of the art. To further the preceding example, the Associated Press® data stream may be filtered by a predetermined geography attribute to generate an enriched data stream.

[0083] The enriched data stream is then subsequently stored, as described in greater detail herein, in a predetermined location. In various embodiments, the enriched data stream is cached by the storage 424 component to provide a local version of the enriched data stream. In certain embodiments, the cached, enriched data stream is implemented to be "replayed" by the cognitive engine 320. In one embodiment, the replaying of the cached, enriched data stream allows incremental ingestion of the enriched data stream instead of ingesting the entire enriched data stream at one time. In various embodiments, one or more enrichment agents 425 are implemented to be invoked by the enrichment component 423 to perform one or more enrichment operations described in greater detail herein.

[0084] In various embodiments, the graph query engine 326 is implemented to receive and process queries such that they can be bridged into a cognitive graph, as described in greater detail herein, through the use of a bridging agent. In these embodiments, the graph query engine 326 may include a query 426 component, a translate 427 component, a bridge 428 component, and one or more bridging agents 429.

[0085] In various embodiments, the query 426 component is implemented to support natural language queries. In these and other embodiments, the query 426 component receives queries, processes them (e.g., using NLP processes), and then maps the processed query to a target cognitive graph. In various embodiments, the translate 427 component is implemented to convert the processed queries provided by the query 426 component into a form that can be used to query a target cognitive graph. To further differentiate the distinction between the functionality respectively provided by the query 426 and translate 427 components, the query 426 component is oriented toward understanding a query from a user. In contrast, the translate 427 component is oriented to translating a query that is understood into a form that can be used to query a cognitive graph.

[0086] In various embodiments, the bridge 428 component is implemented to generate an answer to a query provided by the translate 427 component. In certain embodiments, the bridge 428 component is implemented to provide domain-specific responses when bridging a translated query to a cognitive graph. For example, the same query bridged to a target cognitive graph by the bridge 428 component may result in different answers for different domains, dependent upon domain-specific bridging operations performed by the bridge 428 component.

[0087] To further differentiate the distinction between the translate 427 component and the bridging 428 component, the translate 427 component relates to a general domain transla-

tion of a question. In contrast, the bridging **428** component allows the question to be asked in the context of a specific domain (e.g., healthcare, travel, etc.), given what is known about the data. In certain embodiments, the bridging **428** component is implemented to process what is known about the translated query, in the context of the user, to provide an answer that is relevant to a specific domain.

[0088] As an example, a user may ask, “Where should I eat today?” If the user has been prescribed a particular health regimen, the bridging **428** component may suggest a restaurant with a “heart healthy” menu. However, if the user is a business traveler, the bridging **428** component may suggest the nearest restaurant that has the user’s favorite food. In various embodiments, the bridging **428** component may provide answers, or suggestions, that are composed and ranked according to a specific domain of use. In various embodiments, the bridging agent **429** is implemented to interact with the bridging component **428** to perform bridging operations described in greater detail herein. In these embodiments, the bridging agent interprets a translated query generated by the query **426** component within a predetermined user context, and then maps it to predetermined nodes and links within a target cognitive graph.

[0089] In various embodiments, the insight/learning engine **330** is implemented to encapsulate a predetermined algorithm, which is then applied to a target cognitive graph to generate a result, such as a cognitive insight or a recommendation. In certain embodiments, one or more such algorithms may contribute to answering a specific question and provide additional cognitive insights or recommendations. In these and other embodiments, the insight/learning engine **330** is implemented to perform insight/learning operations, described in greater detail herein. In various embodiments, the insight/learning engine **330** may include a discover/visibility **430** component, a predict **431** component, a rank/recommend **432** component, and one or more insight **433** agents.

[0090] In various embodiments, the discover/visibility **430** component is implemented to provide detailed information related to a predetermined topic, such as a subject or an event, along with associated historical information. In certain embodiments, the predict **431** component is implemented to perform predictive operations to provide insight into what may next occur for a predetermined topic. In various embodiments, the rank/recommend **432** component is implemented to perform ranking and recommendation operations to provide a user prioritized recommendations associated with a provided cognitive insight.

[0091] In certain embodiments, the insight/learning engine **330** may include additional components. For example the additional components may include classification algorithms, clustering algorithms, and so forth. Skilled practitioners of the art will realize that many such additional components are possible and that the foregoing is not intended to limit the spirit, scope or intent of the invention. In various embodiments, the insights agents **433** are implemented to create a visual data story, highlighting user-specific insights, relationships and recommendations. As a result, it can share, operationalize, or track business insights in various embodiments. In various embodiments, the learning agent **434** work in the background to continually update the cognitive graph, as described in greater detail herein, from each unique interaction with data and users.

[0092] In various embodiments, the destination agents **336** are implemented to publish cognitive insights to a consumer of cognitive insight data. Examples of such consumers of cognitive insight data include target databases, business intelligence applications, and mobile applications. In various embodiments, the destination agents **336** may include a Hypertext Transfer Protocol (HTTP) stream **440** agent, an API connectors **441** agent, a databases **442** agent, a message engines **443** agent, a mobile push notification **444** agent, and one or more custom destination **446** agents. Skilled practitioners of the art will realize that other types of destination agents **318** may be used in various embodiments and the foregoing is not intended to limit the spirit, scope or intent of the invention. In certain embodiments, each of the destination agents **318** has a corresponding API.

[0093] In various embodiments, the HTTP stream **440** agent is implemented for providing various HTTP streams of cognitive insight data to a predetermined cognitive data consumer. In these embodiments, the provided HTTP streams may include various HTTP data elements familiar to those of skill in the art. In certain embodiments, the HTTP streams of data are provided in near-real-time. In various embodiments, the API connectors **441** agent is implemented to manage interactions with one or more predetermined APIs that are external to the cognitive platform **310**. As an example, various target databases, business intelligence applications, and mobile applications may each have their own unique API.

[0094] In various embodiments, the databases **442** agent is implemented for provision of cognitive insight data to one or more target databases familiar to those of skill in the art. For example, the target database may include a SQL, NoSQL, delimited flat file, or other form of database. In these embodiments, the provided cognitive insight data may include a single data element, a single data record or file, or a plurality of data records or files. In certain embodiments, the data may be provided to more than one cognitive data consumer and the provided data may be in a homogenous or heterogeneous form. In various embodiments, the message engines **443** agent is implemented to provide cognitive insight data to one or more message engines, such as a message queue (MQ) system, a message bus, a message broker, an enterprise service bus (ESB), and so forth. Skilled practitioners of the art will realize that there are many such examples of message engines with which the message engines **443** agent may interact and the foregoing is not intended to limit the spirit, scope or intent of the invention.

[0095] In various embodiments, the custom destination agents **420**, which are purpose-built, are developed through the use of the development environment **314**, described in greater detail herein. Examples of custom destination agents **420** include destination agents for various electronic medical record (EMR) systems at various healthcare facilities. Such EMR systems typically collect a variety of healthcare information, much of it the same, yet it may be collected, stored and provided in different ways. In this example, the custom destination agents **420** allow such EMR systems to receive cognitive insight data in a form they can use.

[0096] In various embodiments, data that has been cleansed, normalized and enriched by the dataset engine, as described in greater detail herein, is provided by a destination agent **336** to a predetermined destination, likewise described in greater detail herein. In these embodiments, neither the graph query engine **326** nor the insight/learning engine **330** are implemented to perform their respective functions.

[0097] In various embodiments, the foundation components 334 are implemented to enable the dataset engine 322, the graph query engine 326, and the insight/learning engine 330 to perform their respective operations and processes. In these and other embodiments, the foundation components 334 may include an NLP core 436 component, an NLP services 437 component, and a dynamic pipeline engine 438. In various embodiments, the NLP core 436 component is implemented to provide a set of predetermined NLP components for performing various NLP operations described in greater detail herein.

[0098] In these embodiments, certain of these NLP core components are surfaced through the NLP services 437 component, while some are used as libraries. Examples of operations that are performed with such components include dependency parsing, parts-of-speech tagging, sentence pattern detection, and so forth. In various embodiments, the NLP services 437 component is implemented to provide various internal NLP services, which are used to perform entity detection, summarization, and other operations, likewise described in greater detail herein. In these embodiments, the NLP services 437 component is implemented to interact with the NLP core 436 component to provide predetermined NLP services, such as summarizing a target paragraph.

[0099] In various embodiments, the dynamic pipeline engine 438 is implemented to interact with the dataset engine 322 to perform various operations related to receiving one or more sets of data from one or more sourcing agents, apply enrichment to the data, and then provide the enriched data to a predetermined destination. In these and other embodiments, the dynamic pipeline engine 438 manages the distribution of these various operations to a predetermined compute cluster and tracks versioning of the data as it is processed across various distributed computing resources. In certain embodiments, the dynamic pipeline engine 438 is implemented to perform data sovereignty management operations to maintain sovereignty of the data.

[0100] In various embodiments, the platform data 338 includes various data repositories, described in greater detail herein, that are accessed by the cognitive platform 310 to generate cognitive insights. In these embodiments, the platform data 338 repositories may include repositories of dataset metadata 456, cognitive graphs 457, models 459, crawl data 460, and management metadata 461. In various embodiments, the dataset metadata 456 is associated with curated data 458 contained in the repository of cognitive graphs 457. In these and other embodiments, the repository of dataset metadata 456 contains dataset metadata that supports operations performed by the storage 424 component of the dataset engine 322. For example, if a Mongo® NoSQL database with ten million items is being processed, and the cognitive platform 310 fails after ingesting nine million of the items, then the dataset metadata 456 may be able to provide a checkpoint that allows ingestion to continue at the point of failure instead restarting the ingestion process.

[0101] Those of skill in the art will realize that the use of such dataset metadata 456 in various embodiments allows the dataset engine 322 to be stateful. In certain embodiments, the dataset metadata 456 allows support of versioning. For example versioning may be used to track versions of modifications made to data, such as in data enrichment processes described in greater detail herein. As another example, geotagging information may have been applied to a set of data during a first enrichment process, which creates a first version

of enriched data. Adding sentiment data to the same million records during a second enrichment process creates a second version of enriched data. In this example, the dataset metadata stored in the dataset metadata 456 provides tracking of the different versions of the enriched data and the differences between the two.

[0102] In various embodiments, the repository of cognitive graphs 457 is implemented to store cognitive graphs generated, accessed, and updated by the cognitive engine 320 in the process of generating cognitive insights. In various embodiments, the repository of cognitive graphs 457 may include one or more repositories of curated data 458, described in greater detail herein. In certain embodiments, the repositories of curated data 458 includes data that has been curated by one or more users, machine operations, or a combination of the two, by performing various sourcing, filtering, and enriching operations described in greater detail herein. In these and other embodiments, the curated data 458 is ingested by the cognitive platform 310 and then processed, as likewise described in greater detail herein, to generate cognitive insights. In various embodiments, the repository of models 459 is implemented to store models that are generated, accessed, and updated by the cognitive engine 320 in the process of generating cognitive insights. As used herein, models broadly refer to machine learning models. In certain embodiments, the models include one or more statistical models.

[0103] In various embodiments, the crawl framework 452 is implemented to support various crawlers 454 familiar to skilled practitioners of the art. In certain embodiments, the crawlers 454 are custom configured for various target domains. For example, different crawlers 454 may be used for various travel forums, travel blogs, travel news and other travel sites. In various embodiments, data collected by the crawlers 454 is provided by the crawl framework 452 to the repository of crawl data 460. In these embodiments, the collected crawl data is processed and then stored in a normalized form in the repository of crawl data 460. The normalized data is then provided to SQL/NoSQL database 417 agent, which in turn provides it to the dataset engine 322. In one embodiment, the crawl database 460 is a NoSQL database, such as Mongo®.

[0104] In various embodiments, the repository of management metadata 461 is implemented to store user-specific metadata used by the management console 312 to manage accounts (e.g., billing information) and projects. In certain embodiments, the user-specific metadata stored in the repository of management metadata 461 is used by the management console 312 to drive processes and operations within the cognitive platform 310 for a predetermined project. In various embodiments, the user-specific metadata stored in the repository of management metadata 461 is used to enforce data sovereignty. It will be appreciated that many such embodiments are possible and the foregoing is not intended to limit the spirit, scope or intent of the invention.

[0105] Referring now to FIG. 4c, the cloud infrastructure 340 may include a cognitive cloud management 342 component and a cloud analytics infrastructure 344 component in various embodiments. Current examples of a cloud infrastructure 340 include Amazon Web Services (AWS®), available from Amazon.com® of Seattle, Wash., IBM® Softlayer, available from International Business Machines of Armonk, N.Y., and Nebula/Openstack, a joint project between Rackspace Hosting®, of Windcrest, Tex., and the National Aero-

nautics and Space Administration (NASA). In these embodiments, the cognitive cloud management **342** component may include a management playbooks **468** sub-component, a cognitive cloud management console **469** sub-component, a data console **470** sub-component, an asset repository **471** sub-component. In certain embodiments, the cognitive cloud management **342** component may include various other sub-components.

[0106] In various embodiments, the management playbooks **468** sub-component is implemented to automate the creation and management of the cloud analytics infrastructure **344** component along with various other operations and processes related to the cloud infrastructure **340**. As used herein, "management playbooks" broadly refers to any set of instructions or data, such as scripts and configuration data, that is implemented by the management playbooks **468** sub-component to perform its associated operations and processes.

[0107] In various embodiments, the cognitive cloud management console **469** sub-component is implemented to provide a user visibility and management controls related to the cloud analytics infrastructure **344** component along with various other operations and processes related to the cloud infrastructure **340**. In various embodiments, the data console **470** sub-component is implemented to manage platform data **338**, described in greater detail herein. In various embodiments, the asset repository **471** sub-component is implemented to provide access to various cognitive cloud infrastructure assets, such as asset configurations, machine images, and cognitive insight stack configurations.

[0108] In various embodiments, the cloud analytics infrastructure **344** component may include a data grid **472** sub-component, a distributed compute engine **474** sub-component, and a compute cluster management **476** sub-component. In these embodiments, the cloud analytics infrastructure **344** component may also include a distributed object storage **478** sub-component, a distributed full text search **480** sub-component, a document database **482** sub-component, a graph database **484** sub-component, and various other sub-components. In various embodiments, the data grid **472** sub-component is implemented to provide distributed and shared memory that allows the sharing of objects across various data structures. One example of a data grid **472** sub-component is Redis, an open-source, networked, in-memory, key-value data store, with optional durability, written in ANSI C. In various embodiments, the distributed compute engine **474** sub-component is implemented to allow the cognitive platform **310** to perform various cognitive insight operations and processes in a distributed computing environment. Examples of such cognitive insight operations and processes include batch operations and streaming analytics processes.

[0109] In various embodiments, the compute cluster management **476** sub-component is implemented to manage various computing resources as a compute cluster. One such example of such a compute cluster management **476** sub-component is Mesos/Nimbus, a cluster management platform that manages distributed hardware resources into a single pool of resources that can be used by application frameworks to efficiently manage workload distribution for both batch jobs and long-running services. In various embodiments, the distributed object storage **478** sub-component is implemented to manage the physical storage and retrieval of distributed objects (e.g., binary file, image, text, etc.) in a cloud

environment. Examples of a distributed object storage **478** sub-component include Amazon S3®, available from Amazon.com of Seattle, Wash., and Swift, an open source, scalable and redundant storage system.

[0110] In various embodiments, the distributed full text search **480** sub-component is implemented to perform various full text search operations familiar to those of skill in the art within a cloud environment. In various embodiments, the document database **482** sub-component is implemented to manage the physical storage and retrieval of structured data in a cloud environment. Examples of such structured data include social, public, private, and device data, as described in greater detail herein. In certain embodiments, the structured data includes data that is implemented in the JavaScript Object Notation (JSON) format. One example of a document database **482** sub-component is Mongo, an open source cross-platform document-oriented database. In various embodiments, the graph database **484** sub-component is implemented to manage the physical storage and retrieval of cognitive graphs. One example of a graph database **484** sub-component is GraphDB, an open source graph database familiar to those of skill in the art.

[0111] FIG. 5 is a simplified process diagram of cognitive inference and learning system (CILS) operations performed in accordance with an embodiment of the invention. In various embodiments, these CILS operations may include a perceive **506** phase, a relate **508** phase, an operate **510** phase, a process and execute **512** phase, and a learn **514** phase. In these and other embodiments, the CILS **118** shown in FIG. 2 is implemented to mimic cognitive processes associated with the human brain. In various embodiments, the CILS operations are performed through the implementation of a cognitive platform **310**, described in greater detail herein. In these and other embodiments, the cognitive platform **310** may be implemented within a cloud analytics infrastructure **344**, which in turn is implemented within a cloud infrastructure **340**, likewise described in greater detail herein.

[0112] In various embodiments, multi-site, multi-structured source streams **504** are provided by sourcing agents, as described in greater detail herein. In these embodiments, the source streams **504** are dynamically ingested in real-time during the perceive **506** phase, and based upon a predetermined context, extraction, parsing, and tagging operations are performed on language, text and images contained in the source streams **504**. Automatic feature extraction and modeling operations are then performed with the previously processed source streams **504** during the relate **508** phase to generate queries to identify related data (i.e., corpus expansion).

[0113] In various embodiments, operations are performed during the operate **510** phase to discover, summarize and prioritize various concepts, which are in turn used to generate actionable recommendations and notifications associated with predetermined plan-based optimization goals. The resulting actionable recommendations and notifications are then processed during the process and execute **512** phase to provide cognitive insights, such as recommendations, to various predetermined destinations and associated application programming interfaces (APIs) **524**.

[0114] In various embodiments, features from newly-observed data are automatically extracted from user feedback during the learn **514** phase to improve various analytical models. In these embodiments, the learn **514** phase includes feedback on observations generated during the relate **508**

phase, which is provided to the perceive **506** phase. Likewise, feedback on decisions resulting from operations performed during the operate **510** phase, and feedback on results resulting from operations performed during the process and execute **512** phase, are also provided to the perceive **506** phase.

[0115] In various embodiments, user interactions result from operations performed during the process and execute **512** phase. In these embodiments, data associated with the user interactions are provided to the perceive **506** phase as unfolding interactions **522**, which include events that occur external to the CILS operations described in greater detail herein. As an example, a first query from a user may be submitted to the CILS system, which in turn generates a first cognitive insight, which is then provided to the user. In response, the user may respond by providing a first response, or perhaps a second query, either of which is provided in the same context as the first query. The CILS receives the first response or second query, performs various CILS operations, and provides the user a second cognitive insight. As before, the user may respond with a second response or a third query, again in the context of the first query. Once again, the CILS performs various CILS operations and provides the user a third cognitive insight, and so forth. In this example, the provision of cognitive insights to the user, and their various associated responses, results in unfolding interactions **522**, which in turn result in a stateful dialog that evolves over time. Skilled practitioners of the art will likewise realize that such unfolding interactions **522**, occur outside of the CILS operations performed by the cognitive platform **310**.

[0116] FIG. 6 depicts the lifecycle of CILS agents implemented in accordance with an embodiment of the invention to perform CILS operations. In various embodiments, the CILS agents lifecycle **602** may include implementation of a sourcing **318** agent, an enrichment **425** agent, a bridging **429** agent, an insight **433** agent, a destination **336** agent, and a learning **434** agent. In these embodiments, the sourcing **318** agent is implemented to source a variety of multi-site, multi-structured source streams of data described in greater detail herein. These sourced data streams are then provided to an enrichment **425** agent, which then invokes an enrichment component to perform enrichment operations to generate enriched data streams, likewise described in greater detail herein.

[0117] The enriched data streams are then provided to a bridging **429** agent, which is used to perform bridging operations described in greater detail herein. In turn, the results of the bridging operations are provided to an insight **433** agent, which is implemented as described in greater detail herein to create a visual data story, highlighting user-specific insights, relationships and recommendations. The resulting visual data story is then provided to a destination **336** agent, which is implemented to publish cognitive insights to a consumer of cognitive insight data, likewise as described in greater detail herein. In response, the consumer of cognitive insight data provides feedback to a learning **434** agent, which is implemented as described in greater detail herein to provide the feedback to the sourcing agent **318**, at which point the CILS agents lifecycle **602** is continued. From the foregoing, skilled practitioners of the art will recognize that each iteration of the cognitive agents lifecycle **602** provides more informed cognitive insights.

[0118] FIG. 7 is a simplified block diagram of a plurality of cognitive platforms implemented in accordance with an embodiment of the invention within a hybrid cloud infrastruc-

ture. In this embodiment, the hybrid cloud infrastructure **740** includes a cognitive cloud management **342** component, a hosted **704** cognitive cloud environment, and a private **706** network environment. As shown in FIG. 7, the hosted **704** cognitive cloud environment includes a hosted **710** cognitive platform, such as the cognitive platform **310** shown in FIGS. 3, 4a, and 4b. In various embodiments, the hosted **704** cognitive cloud environment may also include a hosted **718** universal knowledge repository and one or more repositories of curated public data **714** and licensed data **716**. Likewise, the hosted **710** cognitive platform may also include a hosted **712** analytics infrastructure, such as the cloud analytics infrastructure **344** shown in FIGS. 3 and 4c.

[0119] As likewise shown in FIG. 7, the private **706** network environment includes a private **720** cognitive platform, such as the cognitive platform **310** shown in FIGS. 3, 4a, and 4b. In various embodiments, the private **706** network cognitive cloud environment may also include a private **728** universal knowledge repository and one or more repositories of application data **724** and private data **726**. Likewise, the private **720** cognitive platform may also include a private **722** analytics infrastructure, such as the cloud analytics infrastructure **344** shown in FIGS. 3 and 4c. In certain embodiments, the private **706** network environment may have one or more private **736** cognitive applications implemented to interact with the private **720** cognitive platform.

[0120] As used herein, a universal knowledge repository broadly refers to a collection of knowledge elements that can be used in various embodiments to generate one or more cognitive insights described in greater detail herein. In various embodiments, these knowledge elements may include facts (e.g., milk is a dairy product), information (e.g., an answer to a question), descriptions (e.g., the color of an automobile), skills (e.g., the ability to install plumbing fixtures), and other classes of knowledge familiar to those of skill in the art. In these embodiments, the knowledge elements may be explicit or implicit. As an example, the fact that water freezes at zero degrees centigrade would be an explicit knowledge element, while the fact that an automobile mechanic knows how to repair an automobile would be an implicit knowledge element.

[0121] In certain embodiments, the knowledge elements within a universal knowledge repository may also include statements, assertions, beliefs, perceptions, preferences, sentiments, attitudes or opinions associated with a person or a group. As an example, user 'A' may prefer the pizza served by a first restaurant, while user 'B' may prefer the pizza served by a second restaurant. Furthermore, both user 'A' and 'B' are firmly of the opinion that the first and second restaurants respectively serve the very best pizza available. In this example, the respective preferences and opinions of users 'A' and 'B' regarding the first and second restaurant may be included in the universal knowledge repository **880** as they are not contradictory. Instead, they are simply knowledge elements respectively associated with the two users and can be used in various embodiments for the generation of various cognitive insights, as described in greater detail herein.

[0122] In various embodiments, individual knowledge elements respectively associated with the hosted **718** and private **728** universal knowledge repositories may be distributed. In one embodiment, the distributed knowledge elements may be stored in a plurality of data stores familiar to skilled practitioners of the art. In this embodiment, the distributed knowledge elements may be logically unified for various implemen-

tations of the hosted **718** and private **728** universal knowledge repositories. In certain embodiments, the hosted **718** and private **728** universal knowledge repositories may be respectively implemented in the form of a hosted or private universal cognitive graph. In these embodiments, nodes within the hosted or private universal graph contain one or more knowledge elements.

[0123] In various embodiments, a secure tunnel **730**, such as a virtual private network (VPN) tunnel, is implemented to allow the hosted **710** cognitive platform and the private **720** cognitive platform to communicate with one another. In these various embodiments, the ability to communicate with one another allows the hosted **710** and private **720** cognitive platforms to work collaboratively when generating cognitive insights described in greater detail herein. In various embodiments, the hosted **710** cognitive platform accesses knowledge elements stored in the hosted **718** universal knowledge repository and data stored in the repositories of curated public data **714** and licensed data **716** to generate various cognitive insights. In certain embodiments, the resulting cognitive insights are then provided to the private **720** cognitive platform, which in turn provides them to the one or more private cognitive applications **736**.

[0124] In various embodiments, the private **720** cognitive platform accesses knowledge elements stored in the private **728** universal knowledge repository and data stored in the repositories of application data **724** and private data **726** to generate various cognitive insights. In turn, the resulting cognitive insights are then provided to the one or more private cognitive applications **736**. In certain embodiments, the private **720** cognitive platform accesses knowledge elements stored in the hosted **718** and private **728** universal knowledge repositories and data stored in the repositories of curated public data **714**, licensed data **716**, application data **724** and private data **726** to generate various cognitive insights. In these embodiments, the resulting cognitive insights are in turn provided to the one or more private cognitive applications **736**.

[0125] In various embodiments, the secure tunnel **730** is implemented for the hosted **710** cognitive platform to provide **732** predetermined data and knowledge elements to the private **720** cognitive platform. In one embodiment, the provision **732** of predetermined knowledge elements allows the hosted **718** universal knowledge repository to be replicated as the private **728** universal knowledge repository. In another embodiment, the provision **732** of predetermined knowledge elements allows the hosted **718** universal knowledge repository to provide updates **734** to the private **728** universal knowledge repository. In certain embodiments, the updates **734** to the private **728** universal knowledge repository do not overwrite other data. Instead, the updates **734** are simply added to the private **728** universal knowledge repository.

[0126] In one embodiment, knowledge elements that are added to the private **728** universal knowledge repository are not provided to the hosted **718** universal knowledge repository. As an example, an airline may not wish to share private information related to its customer's flights, the price paid for tickets, their awards program status, and so forth. In another embodiment, predetermined knowledge elements that are added to the private **728** universal knowledge repository may be provided to the hosted **718** universal knowledge repository. As an example, the operator of the private **720** cognitive platform may decide to license predetermined knowledge elements stored in the private **728** universal knowledge

repository to the operator of the hosted **710** cognitive platform. To continue the example, certain knowledge elements stored in the private **728** universal knowledge repository may be anonymized prior to being provided for inclusion in the hosted **718** universal knowledge repository. In one embodiment, only private knowledge elements are stored in the private **728** universal knowledge repository. In this embodiment, the private **720** cognitive platform may use knowledge elements stored in both the hosted **718** and private **728** universal knowledge repositories to generate cognitive insights. Skilled practitioners of the art will recognize that many such embodiments are possible and the foregoing is not intended to limit the spirit, scope or intent of the invention.

[0127] FIG. **8** is a simplified process flow diagram of composite cognitive insight generation and feedback operations performed in accordance with an embodiment of the invention. As used herein, a composite cognitive insight broadly refers to a set of cognitive insights generated as a result of orchestrating a predetermined set of independent cognitive agents, referred to herein as insight agents. In various embodiments, the insight agents use a cognitive graph, such as an application cognitive graph **882** as their data source to respectively generate individual cognitive insights. As used herein, an application cognitive graph **882** broadly refers to a cognitive graph that is associated with a predetermined cognitive application **304**. In certain embodiments, different cognitive applications **304** may interact with different application cognitive graphs **882** to generate individual cognitive insights for a user. In various embodiments, the resulting individual cognitive insights are then composed to generate a set of cognitive composite insights, which in turn is provided to a user in the form of a cognitive insight summary **848**.

[0128] In various embodiments, the orchestration of the selected insight agents is performed by the cognitive insight/learning engine **330** shown in FIGS. **3** and **4a**. In certain embodiments, a predetermined subset of insight agents is selected to provide composite cognitive insights to satisfy a graph query **844**, a contextual situation, or some combination thereof. For example, it may be determined, as described in greater detail herein, that a particular subset of insight agents may be suited to provide a composite cognitive insight related to a particular user of a particular device, at a particular location, at a particular time, for a particular purpose.

[0129] In certain embodiments, the insight agents are selected for orchestration as a result of receiving direct or indirect input data **842** from a user. In various embodiments, the direct user input may be a natural language inquiry. In certain embodiments, the indirect user input data **842** may include the location of a user's device or the purpose for which it is being used. As an example, the Geographical Positioning System (GPS) coordinates of the location of a user's mobile device may be received as indirect user input data **842**. As another example, a user may be using the integrated camera of their mobile device to take a photograph of a location, such as a restaurant, or an item, such as a food product. In certain embodiments, the direct or indirect user input data **842** may include personal information that can be used to identify the user. Skilled practitioners of the art will recognize that many such embodiments are possible and the foregoing is not intended to limit the spirit, scope or intent of the invention.

[0130] In various embodiments, composite cognitive insight generation and feedback operations may be performed in various phases. In this embodiment, these phases

include a data lifecycle **840** phase, a learning **838** phase, and an application/insight composition **840** phase. In the data lifecycle **836** phase, a predetermined instantiation of a cognitive platform **810** sources social data **812**, public data **814**, licensed data **816**, and proprietary data **818** from various sources as described in greater detail herein. In various embodiments, an example of a cognitive platform **810** instantiation is the cognitive platform **310** shown in FIGS. **3**, **4a**, and **4b**. In this embodiment, the instantiation of a cognitive platform **810** includes a source **806** component, a process **808** component, a deliver **810** component, a cleanse **820** component, an enrich **822** component, a filter/transform **824** component, and a repair/reject **826** component. Likewise, as shown in FIG. **8**, the process **808** component includes a repository of models **828**, described in greater detail herein.

[0131] In various embodiments, the process **808** component is implemented to perform various composite insight generation and other processing operations described in greater detail herein. In these embodiments, the process **808** component is implemented to interact with the source **806** component, which in turn is implemented to perform various data sourcing operations described in greater detail herein. In various embodiments, the sourcing operations are performed by one or more sourcing agents, as likewise described in greater detail herein. The resulting sourced data is then provided to the process **808** component. In turn, the process **808** component is implemented to interact with the cleanse **820** component, which is implemented to perform various data cleansing operations familiar to those of skill in the art. As an example, the cleanse **820** component may perform data normalization or pruning operations, likewise known to skilled practitioners of the art. In certain embodiments, the cleanse **820** component may be implemented to interact with the repair/reject **826** component, which in turn is implemented to perform various data repair or data rejection operations known to those of skill in the art.

[0132] Once data cleansing, repair and rejection operations are completed, the process **808** component is implemented to interact with the enrich **822** component, which is implemented in various embodiments to perform various data enrichment operations described in greater detail herein. Once data enrichment operations have been completed, the process **808** component is likewise implemented to interact with the filter/transform **824** component, which in turn is implemented to perform data filtering and transformation operations described in greater detail herein.

[0133] In various embodiments, the process **808** component is implemented to generate various models, described in greater detail herein, which are stored in the repository of models **828**. The process **808** component is likewise implemented in various embodiments to use the sourced data to generate one or more cognitive graphs, such as an application cognitive graph **882**, as described in greater detail herein. In various embodiments, the process **808** component is implemented to gain an understanding of the data sourced from the sources of social data **812**, public data **814**, licensed data **816**, and proprietary data **818**, which assist in the automated generation of the application cognitive graph **882**.

[0134] The process **808** component is likewise implemented in various embodiments to perform bridging **846** operations, described in greater detail herein, to access the application cognitive graph **882**. In certain embodiments, the bridging **846** operations are performed by bridging agents, likewise described in greater detail herein. In various embodi-

ments, the application cognitive graph **882** is accessed by the process **808** component during the learning **836** phase of the composite cognitive insight generation operations.

[0135] In various embodiments, a cognitive application **304** is implemented to receive input data **842** associated with an individual user or a group of users. In these embodiments, the input data **842** may be direct, such as a user query or mouse click, or indirect, such as the current time or Geographical Positioning System (GPS) data received from a mobile device associated with a user. In various embodiments, the indirect input data **842** may include contextual data, described in greater detail herein. Once it is received, the input data **842** is then submitted by the cognitive application **304** to a graph query engine **326** during the application/insight composition **840** phase. In turn, the graph query engine **326** processes the submitted input data **842** to generate a graph query **844**, as described in greater detail herein. The graph query **844** is then used to query the application cognitive graph **882**, which results in the generation of one or more composite cognitive insights, likewise described in greater detail herein. In certain embodiments, the graph query **844** uses predetermined knowledge elements stored in the universal knowledge repository **880** when querying the application cognitive graph **882** to generate the one or more composite insights.

[0136] In various embodiments, a set of contextually-related interactions between a cognitive application **304** and the application cognitive graph **882** are represented as a corresponding set of nodes in a predetermined cognitive session graph, which is then stored in a repository of cognitive session graphs '1' through 'n' **852**. As used herein, a cognitive session graph broadly refers to a cognitive graph whose nodes are associated with a cognitive session. As used herein, a cognitive session broadly refers to a predetermined user, group of users, theme, topic, issue, question, intent, goal, objective, task, assignment, process, situation, requirement, condition, responsibility, location, period of time, or any combination thereof.

[0137] As an example, the application cognitive graph **882** may be unaware of a particular user's preferences, which are likely stored in a corresponding user profile. To further the example, a user may typically choose a particular brand or manufacturer when shopping for a given type of product, such as cookware. A record of each query regarding that brand of cookware, or its selection, is iteratively stored in a predetermined cognitive session graph that is associated with the user and stored in a repository of cognitive session graphs '1' through 'n' **852**. As a result, the preference of that brand of cookware is ranked higher, and is presented in response to contextually-related queries, even when the preferred brand of cookware is not explicitly referenced by the user. To continue the example, the user may make a number of queries over a period of days or weeks, yet the queries are all associated with the same cognitive session graph that is associated with the user and stored in a repository of cognitive session graphs '1' through 'n' **852**, regardless of when each query is made.

[0138] As another example, a user queries a cognitive application **304** during business hours to locate an upscale restaurant located close their place of business. As a result, a first cognitive session graph stored in a repository of cognitive session graphs '1' through 'n' **852** is associated with the user's query, which results in the provision of composite cognitive insights related to restaurants suitable for business

meetings. To continue the example, the same user queries the same cognitive application **304** during the weekend to locate a casual restaurant located close to their home. As a result, a second cognitive session graph stored in a repository of cognitive session graphs '1' through 'n' **852** is associated with the user's query, which results in the provision of composite cognitive insights related to restaurants suitable for family meals. In these examples, the first and second cognitive session graphs are both associated with the same user, but for two different purposes, which results in the provision of two different sets of composite cognitive insights.

[0139] As yet another example, a group of customer support representatives is tasked with resolving technical issues customers may have with a predetermined product. In this example, the product and the group of customer support representatives are collectively associated with a predetermined cognitive session graph stored in a repository of cognitive session graphs '1' through 'n' **852**. To continue the example, individual customer support representatives may submit queries related to the product to a cognitive application **304**, such as a knowledge base application. In response, a predetermined cognitive session graph stored in a repository of cognitive session graphs '1' through 'n' **852** is used, along with the universal knowledge repository **880** and application cognitive graph **882**, to generate individual or composite cognitive insights to resolve a technical issue for a customer. In this example, the cognitive application **304** may be queried by the individual customer support representatives at different times during some predetermined time interval, yet the same cognitive session graph stored in a repository of cognitive session graphs '1' through 'n' **852** is used to generate composite cognitive insights related to the product.

[0140] In various embodiments, each cognitive session graph associated with a user and stored in a repository of cognitive session graphs '1' through 'n' **852** includes one or more direct or indirect user queries represented as nodes, and the time at which they were asked, which are in turn linked **854** to nodes that appear in the application cognitive graph **882**. In certain embodiments, each individual session graph that is associated with the user and stored in a repository of cognitive session graphs '1' through 'n' **852** introduces edges that are not already present in the application cognitive graph **882**. More specifically, each of the session graphs that is associated with the user and stored in a repository of cognitive session graphs '1' through 'n' **852** establishes various relationships that the application cognitive graph **882** does not already have.

[0141] In various embodiments, individual graph queries **844** associated with a predetermined session graph stored in a repository of cognitive session graphs '1' through 'n' **852** are likewise provided to predetermined insight agents to perform various kinds of analyses. In certain embodiments, each insight agent performs a different kind of analysis. In various embodiments, different insight agents may perform the same, or similar, analyses. In certain embodiments, different agents performing the same or similar analyses may be competing between themselves.

[0142] For example, a user may be a realtor that has a young, upper middle-class, urban-oriented clientele that typically enjoys eating at trendy restaurants that are in walking distance of where they live. As a result, the realtor may be interested in knowing about new or popular restaurants that are in walking distance of their property listings that have a young, middle-class clientele. In this example, the user's

queries may result the assignment of predetermined insight agents to perform analysis of various social media interactions to identify such restaurants that have received favorable reviews. To continue the example, the resulting composite insights may be provided as a ranked list of candidate restaurants that may be suitable venues for the realtor to meet his clients.

[0143] In various embodiments, the process **808** component is implemented to provide these composite cognitive insights to the deliver **810** component, which in turn is implemented to deliver the composite cognitive insights in the form of a cognitive insight summary **848** to the cognitive application **304**. In these embodiments, the cognitive platform **810** is implemented to interact with an insight front-end **856** component, which provides a composite insight and feedback interface with the cognitive application **304**. In certain embodiments, the insight front-end **856** component includes an insight Application Program Interface (API) **858** and a feedback API **860**, described in greater detail herein. In these embodiments, the insight API **858** is implemented to convey the cognitive insight summary **848** to the cognitive application **304**. Likewise, the feedback API **860** is used to convey associated direct or indirect user feedback **862** to the cognitive platform **810**. In certain embodiments, the feedback API **860** provides the direct or indirect user feedback **862** to the repository of models **828** described in greater detail herein.

[0144] To continue the preceding example, the user may have received a list of candidate restaurants that may be suitable venues for meeting his clients. However, one of his clients has a pet that they like to take with them wherever they go. As a result, the user provides feedback **862** that he is looking for a restaurant that is pet-friendly. The provided feedback **862** is in turn provided to the insight agents to identify candidate restaurants that are also pet-friendly. In this example, the feedback **862** is stored in the appropriate cognitive session graph stored in a repository of cognitive session graphs '1' through 'n' **852** associated with the user and their original query.

[0145] In various embodiments, as described in the descriptive text associated with FIG. 5, learning operations are iteratively performed during the learning **838** phase to provide more accurate and useful composite cognitive insights. In certain of these embodiments, feedback **862** received from the user is stored in a predetermined cognitive session graph that is associated with the user and stored in a repository of cognitive session graphs '1' through 'n' **852**, which is then used to provide more accurate composite cognitive insights in response to subsequent contextually-relevant queries from the user.

[0146] As an example, composite cognitive insights provided by a particular insight agent related to a first subject may not be relevant or particularly useful to a user of the cognitive application **304**. As a result, the user provides feedback **862** to that effect, which in turn is stored in the appropriate cognitive session graph that is associated with the user and stored in a repository of cognitive session graphs '1' through 'n' **852**. Accordingly, subsequent insights provided by the insight agent related the first subject may be ranked lower, or not provided, within a cognitive insight summary **848** to the user. Conversely, the same insight agent may provide excellent insights related to a second subject, resulting in positive feedback **862** being received from the user. The positive feedback **862** is likewise stored in the appropriate cognitive session graph that is associated with the user and stored in

a repository of cognitive session graphs '1' through 'n' 852. As a result, subsequent insights provided by the insight agent related to the second subject may be ranked higher within a cognitive insight summary 848 provided to the user.

[0147] In various embodiments, the composite insights provided in each cognitive insight summary 848 to the cognitive application 304, and corresponding feedback 862 received from a user in return, is provided in the form of one or more insight streams 864 to an associated cognitive session graph stored in a repository of cognitive session graphs '1' through 'n' 852. In these and other embodiments, the insight streams 864 may contain information related to the user of the cognitive application 304, the time and date of the provided composite cognitive insights and related feedback 862, the location of the user, and the device used by the user.

[0148] As an example, a query related to upcoming activities that is received at 10:00 AM on a Saturday morning from a user's home may return composite cognitive insights related to entertainment performances scheduled for the weekend. Conversely, the same query received at the same time on a Monday morning from a user's office may return composite insights related to business functions scheduled during the work week. In various embodiments, the information contained in the insight streams 864 is used to rank the composite cognitive insights provided in the cognitive insight summary 848. In certain embodiments, the composite cognitive insights are continually re-ranked as additional insight streams 864 are received. Skilled practitioners of the art will recognize that many such embodiments are possible and the foregoing is not intended to limit the spirit, scope or intent of the invention.

[0149] FIG. 9 is a generalized flowchart of composite cognitive insight and feedback operations performed in accordance with an embodiment of the invention. In this embodiment, composite cognitive insight and feedback operations are begun in step 902, followed by a cognitive application requesting an Application Programming Interface (API) key from a cognitive platform, described in greater detail herein, in step 906. The method by which the API key is requested, generated and provided to the cognitive application is a matter of design choice.

[0150] A cognitive session token is then issued to the cognitive application in step 906, followed by a cognitive session token being returned to the cognitive application in step 908 to establish a composite cognitive insight session. As used herein, a composite cognitive insight session broadly refers to a session with a cognitive application, described in greater detail herein, where composite cognitive insights are generated and presented to a user. In various embodiments, the composite cognitive insight session may also include the receipt of feedback from the user, described in greater detail herein. In one embodiment, the cognitive session token is used to establish a composite cognitive insight session that generates a new cognitive session graph. In another embodiment, the cognitive session token is used to establish a composite cognitive insight session that appends composite cognitive insights and user feedback to an existing cognitive session graph associated with the user.

[0151] In various embodiments, the cognitive session token enables the cognitive application to interact with a cognitive session graph associated with the cognitive session token. In these embodiments, the composite cognitive insight session is perpetuated. For example, a given composite cognitive insight session may last months or even years. In certain

embodiments, the cognitive session token expires after a predetermined period of time. In these embodiments, the cognitive session token is no longer valid once the predetermined period of time expires. The method by which the period of time is determined, and monitored, is a matter of design choice.

[0152] Contextually-relevant composite cognitive insights are then generated and presented, as described in greater detail herein, to the user in step 910. In various embodiments, composite cognitive insights presented to the user are stored in the cognitive session graph associated with the cognitive session token. A determination is then made in step 912 if feedback is received from the user. If not, then a determination is made in step 924 whether to end composite cognitive insights and feedback operations. If not, then a determination is made in step 926 whether the cognitive session token for the target session has expired. If not, then the process is continued, proceeding with step 910. Otherwise, the process is continued, proceeding with step 904. However, if it is determined in step 924 not to end composite cognitive insights and feedback operations, then composite cognitive insights and feedback operations are ended in step 928.

[0153] However, if it was determined in step 912 that feedback was received from the user then the cognitive application provides the feedback, as described in greater detail herein, to the cognitive platform in step 914. A determination is then made in step 916 whether to use the provided feedback to generate contextually-relevant questions for provision to the user. If not, then the process is continued, proceeding with step 924. Otherwise, the cognitive platform uses the provided feedback in step 918 to generate contextually-relevant questions. Then, in step 922, the cognitive platform provides the contextually-relevant questions, along with additional composite insights, to the cognitive application. In turn, the cognitive application provides the contextually-relevant questions and additional composite cognitive insights to the user in step 922 and the process is continued, proceeding with step 912.

[0154] Although the present invention has been described in detail, it should be understood that various changes, substitutions and alterations can be made hereto without departing from the spirit and scope of the invention as defined by the appended claims.

What is claimed is:

1. A system comprising:

- a processor;
- a data bus coupled to the processor; and
- a non-transitory, computer-readable storage medium embodying computer program code, the non-transitory, computer-readable storage medium being coupled to the data bus, the computer program code interacting with a plurality of computer operations and comprising instructions executable by the processor and configured for:
 - receiving streams of data from a plurality of data sources;
 - processing the streams of data from the plurality of data sources, the processing the streams of data from the plurality of data sources performing data enriching to provide data enriched data streams; and,
 - storing the data streams and the data enriched data streams within the universal knowledge repository as a collection of knowledge elements.

2. The system of claim 1, wherein the instructions executable by the processor further comprise instructions for:

generating a cognitive insight based upon the collection of knowledge elements stored within the universal knowledge repository.

3. The system of claim 1, wherein:

the data enriching comprises identifying at least some knowledge elements within the collection of knowledge elements as at least one of facts, opinions, descriptions, and skills.

4. The system of claim 1, wherein:

the data enriching comprises identifying at least some knowledge elements within the collection of knowledge elements as at least one of statements, assertions, beliefs, perceptions, preferences, sentiments, attitudes and opinions and associating identified at least some knowledge elements with an entity responsible for generating the at least one of statements, assertions, beliefs, perceptions, preferences, sentiments, attitudes and opinions.

5. The system of claim 1, wherein:

the universal knowledge repository comprises a private knowledge repository; and,
at least some of the plurality of data streams are private data streams.

6. The system of claim 1, wherein:

the universal knowledge repository comprises a universal cognitive graph.

7. A non-transitory, computer-readable storage medium embodying computer program code, the computer program code comprising computer executable instructions configured for:

receiving streams of data from a plurality of data sources; processing the streams of data from the plurality of data sources, the processing the streams of data from the plurality of data sources performing data enriching to provide data enriched data streams; and,

storing the data streams and the data enriched data streams within the universal knowledge repository as a collection of knowledge elements.

8. The non-transitory, computer-readable storage medium of claim 7, wherein the instructions executable by the processor further comprise instructions for: generating a cognitive insight based upon the collection of knowledge elements stored within the universal knowledge repository.

9. The non-transitory, computer-readable storage medium of claim 7, wherein:

the data enriching comprises identifying at least some knowledge elements within the collection of knowledge elements as at least one of facts, opinions, descriptions, and skills.

10. The non-transitory, computer-readable storage medium of claim 7, wherein:

the data enriching comprises identifying at least some knowledge elements within the collection of knowledge elements as at least one of statements, assertions, beliefs, perceptions, preferences, sentiments, attitudes and opinions and associating identified at least some knowledge elements with an entity responsible for generating the at least one of statements, assertions, beliefs, perceptions, preferences, sentiments, attitudes and opinions.

11. The non-transitory, computer-readable storage medium of claim 7, wherein:

the universal knowledge repository comprises a private knowledge repository; and,
at least some of the plurality of data streams are private data streams.

12. The non-transitory, computer-readable storage medium of claim 7, wherein the universal knowledge repository comprises a universal cognitive graph.

13. The non-transitory, computer-readable storage medium of claim 7, wherein the computer executable instructions are deployable to a client system from a server system at a remote location.

14. The non-transitory, computer-readable storage medium of claim 7, wherein the computer executable instructions are provided by a service provider to a user on an on-demand basis.

* * * * *