



- (51) International Patent Classification: 62/378,152 22 August 2016 (22.08.2016) US
 G06N 5/02 (2006.01) G06F 9/50 (2006.01)
 G06F 9/54 (2006.01)
- (21) International Application Number: PCT/US2017/048044
- (22) International Filing Date: 22 August 2017 (22.08.2017)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:

62/378,143	22 August 2016 (22.08.2016)	US
62/378,146	22 August 2016 (22.08.2016)	US
62/378,147	22 August 2016 (22.08.2016)	US
62/378,150	22 August 2016 (22.08.2016)	US
62/378,151	22 August 2016 (22.08.2016)	US
- (71) Applicant: ORACLE INTERNATIONAL CORPORATION [US/US]; 500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US).
- (72) Inventors: SEETHARAMAN, Ganesh; 500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US). STOJANOVIC, Alexander Sasha; 14 W. Central Avenue, Los Gatos, California 95030 (US). NAMARVAR, Hassan Heidari; 500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US). ALLAN, David; 500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US).
- (74) Agent: MEYER, Sheldon, R. et al.; Tucker Ellis LLP, One Market Plaza, Steuart Tower, Suite 700, San Francisco, California 94105 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

(54) Title: SYSTEM AND METHOD FOR INFERENCING OF DATA TRANSFORMATIONS THROUGH PATTERN DECOMPOSITION




FIGURE 40

(57) Abstract: In accordance with various embodiments, described herein is a system (Data Artificial Intelligence system, Data AI system), for use with a data integration or other computing environment, that leverages machine learning (ML, DataFlow Machine Learning, DFML), for use in managing a flow of data (dataflow, DF), and building complex dataflow software applications (dataflow applications, pipelines). In accordance with an embodiment, the system can provide a service to recommend actions and transformations, on an input data, based on patterns identified from the functional decomposition of a data flow for a software application, including determining possible transformations of the data flow in subsequent applications. Data flows can be decomposed into a model describing transformations of data, predicates, and business rules applied to the data, and attributes used in the data flows.



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

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

Background:

[0003] Many of today's computing environments require an ability to share large amounts of data among different types of software applications. However, distributed applications may be markedly different in their configurations due to, for example, differences in their types of data supported; or their execution environments. An application's configuration may depend, for example, on its application programming interfaces, runtime environment, deployment scheme, lifecycle management, or security management.

[0004] Software design tools intended for use in developing such distributed applications tend to be resource-intensive, often requiring the services of a human domain model expert to curate application and data integrations. Consequently, application developers who are faced with the task of building complex, scalable, distributed applications, that will be used to integrate different types of data among different types of execution environments, must generally expend a substantial amount of manual effort to design, build, and configure those applications.

Summary:

[0005] In accordance with various embodiments, described herein is a system (Data Artificial Intelligence system, Data AI system), for use with a data integration or other computing environment, that leverages machine learning (ML, DataFlow Machine Learning, DFML), for use in managing a flow of data (dataflow, DF), and building complex dataflow software applications (dataflow applications, pipelines). In accordance with an embodiment, the system can provide a service to recommend actions and transformations, on an input data, based on patterns identified from the functional decomposition of a data flow for a software application, including determining possible transformations of the data flow in subsequent applications. Data flows can be decomposed into a model describing transformations of data, predicates, and business rules applied to the data, and attributes used in the data flows.

Brief Description of the Figures:

[0006] **Figure 1** illustrates a system for providing data flow artificial intelligence, in accordance with an embodiment.

[0007] **Figure 2** illustrates an event-driven architecture including an event coordinator for use with a system, in accordance with an embodiment.

[0008] **Figure 3** illustrates the steps in a data flow, in accordance with an embodiment.

[0009] **Figure 4** illustrates an example of a data flow that includes multiple sources, in accordance with an embodiment.

[00010] **Figure 5** illustrates an example use of a data flow with a pipeline, in accordance

with an embodiment.

[00011] Figure 6 illustrates an example use of an ingest/publish engine and ingest/publish service with a pipeline, in accordance with an embodiment.

5 **[00012]** Figure 7 illustrates the process of ingest and training from a HUB, in accordance with an embodiment.

[00013] Figure 8 illustrates the process of building models, in accordance with an embodiment.

[00014] Figure 9 illustrates the process of classifying datasets or entities from newly added HUBs, in accordance with an embodiment.

10 **[00015]** Figure 10 further illustrates the process of classifying datasets or entities, from newly added HUBs, in accordance with an embodiment.

[00016] Figure 11 further illustrates the process of classifying datasets or entities, from newly added HUBs, in accordance with an embodiment.

15 **[00017]** Figure 12 illustrates an object diagram for use in functional type classification, in accordance with an embodiment.

[00018] Figure 13 illustrates an example of a dimension functional type classification, in accordance with an embodiment.

[00019] Figure 14 illustrates an example of a cube functional type classification, in accordance with an embodiment.

20 **[00020]** Figure 15 illustrates an example usage of functional type classification to evaluate a business entity's functional type, in accordance with an embodiment.

[00021] Figure 16 illustrates an object diagram for use in functional transformation, in accordance with an embodiment.

25 **[00022]** Figure 17 illustrates the operation of a recommendation engine, in accordance with an embodiment.

[00023] Figure 18 illustrates the use of a data lake, in accordance with an embodiment.

[00024] Figure 19 illustrates the use of a data-driven strategy to manage a data lake, in accordance with an embodiment.

30 **[00025]** Figure 20 illustrates the use of a process-driven strategy to manage a data lake, in accordance with an embodiment.

[00026] Figure 21 illustrates the use of a pipeline compiler, in accordance with an embodiment.

[00027] Figure 22 illustrates an example pipeline graph, in accordance with an embodiment.

35 **[00028]** Figure 23 illustrates an example of a data pipeline, in accordance with an embodiment.

[00029] Figure 24 illustrates another example of a data pipeline, in accordance with an embodiment.

[00030] Figure 25 illustrates an example of an orchestration pipeline, in accordance with an embodiment.

5 [00031] Figure 26 further illustrates an example of an orchestration pipeline, in accordance with an embodiment.

[00032] Figure 27 illustrates the use of a coordination fabric including a messaging system, in accordance with an embodiment.

10 [00033] Figure 28 further illustrates the use of a coordination fabric including a messaging system, in accordance with an embodiment.

[00034] Figure 29 illustrates an on-premise agent for use with a system, in accordance with an embodiment.

[00035] Figure 30 illustrates a data flow process, in accordance with an embodiment.

15 [00036] Figure 31 illustrates an automated mapping of data types, in accordance with an embodiment.

[00037] Figure 32 illustrates an auto-map service for generation of mappings, in accordance with an embodiment.

[00038] Figure 33 illustrates an example of mapping between a source schema and target schema, in accordance with an embodiment.

20 [00039] Figure 34 illustrates another example of mapping between a source schema and target schema, in accordance with an embodiment.

[00040] Figure 35 illustrates a process for providing automated mapping of data types, in accordance with an embodiment.

25 [00041] Figure 36 illustrates a system that displays one or more semantic actions enabled for accessed data, in accordance with an embodiment.

[00042] Figure 37 illustrates a graphical user interface that displays one or more semantic actions enabled for accessed data, in accordance with an embodiment.

[00043] Figure 38 further illustrates a graphical user interface that displays one or more semantic actions enabled for accessed data, in accordance with an embodiment.

30 [00044] Figure 39 illustrates a process for displaying one or more semantic actions enabled for accessed data, in accordance with an embodiment.

[00045] Figure 40 illustrates a means of identifying a pattern of transformation in a data flow, for one or more functional expressions generated for each of one or more applications, in accordance with an embodiment.

35 [00046] Figure 41 illustrates an example of identifying a pattern of transformation in a data flow, for one or more functional expressions, in accordance with an embodiment.

[00047] Figure 42 illustrates an object diagram for use in identifying a pattern of transformation in a data flow, for one or more functional expressions generated for each of one or more applications, in accordance with an embodiment.

5 [00048] Figure 43 illustrates a process of identifying a pattern of transformation in a data flow, for one or more functional expressions generated for each of one or more applications, in accordance with an embodiment.

[00049] Figure 44 illustrates a system for generating functional type rules, in accordance with an embodiment.

10 [00050] Figure 45 further illustrates a system for generating functional type rules, in accordance with an embodiment.

[00051] Figure 46 illustrates an object diagram for use in generating functional type rules, in accordance with an embodiment.

[00052] Figure 47 illustrates a process for generating a functional type system based on generated one or more rules, in accordance with an embodiment.

15 [00053] Figure 48 illustrates a system for identifying a pattern for use in providing a recommendation for a data flow, based on information provided via a foreign function interface, in accordance with an embodiment.

[00054] Figure 49 further illustrates identifying a pattern for use in providing a recommendation for a data flow, based on information provided via a foreign function interface,

20 in accordance with an embodiment.

[00055] Figure 50 further illustrates identifying a pattern for use in providing a recommendation for a data flow, based on information provided via a foreign function interface, in accordance with an embodiment.

25 [00056] Figure 51 illustrates a process for identifying a pattern for use in providing a recommendation for a data flow, based on information provided via a foreign function interface, in accordance with an embodiment.

[00057] Figure 52 illustrates managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment.

30 [00058] Figure 53 further illustrates managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment.

[00059] Figure 54 further illustrates managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment.

[00060] Figure 55 further illustrates managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment.

35 [00061] Figure 56 further illustrates managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment.

[00062] Figure 57 further illustrates managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment.

[00063] Figure 58 further illustrates managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment.

5 [00064] Figure 59 illustrates a process for managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment.

Detailed Description:

[00065] The foregoing, together with additional embodiments and features thereof will
10 become apparent upon referring to the following description including specification, claims, and accompanying drawings. In the following description, for purposes of explanation, specific details are set forth in order to provide a thorough understanding of various embodiments of the invention. However, it will be apparent that various embodiments can be practiced without these specific details. The following description including specification, claims, and
15 accompanying drawings are not intended to be restrictive.

Introduction

[00066] In accordance with various embodiments, described herein is a system (Data Artificial Intelligence system, Data AI system), for use with a data integration or other
20 computing environment, that leverages machine learning (ML, DataFlow Machine Learning, DFML), for use in managing a flow of data (dataflow, DF), and building complex dataflow software applications (dataflow applications, pipelines).

[00067] In accordance with an embodiment, the system can provide support for auto-mapping of complex data structures, datasets or entities, between one or more sources or
25 targets of data, referred to herein in some embodiments as HUBs. The auto-mapping can be driven by a metadata, schema, and statistical profiling of a dataset; and used to map a source dataset or entity associated with an input HUB, to a target dataset or entity or vice versa, to produce an output data prepared in a format or organization (projection) for use with one or more output HUBs.

30 [00068] In accordance with an embodiment, the system can include a software development component and graphical user interface, referred to herein in some embodiments as a pipeline editor, or Lambda Studio IDE, that provides a visual environment for use with the system, including providing real-time recommendations for performing semantic actions on data accessed from an input HUB, based on an understanding of the meaning or semantics
35 associated with the data.

[00069] In accordance with an embodiment, the system can provide a service to

recommend actions and transformations, on an input data, based on patterns identified from the functional decomposition of a data flow for a software application, including determining possible transformations of the data flow in subsequent applications. Data flows can be decomposed into a model describing transformations of data, predicates, and business rules applied to the data, and attributes used in the data flows.

[00070] In accordance with an embodiment, the system can perform an ontology analysis of a schema definition, to determine the types of data, and datasets or entities, associated with that schema; and generate, or update, a model from a reference schema that includes an ontology defined based on relationships between datasets or entities, and their attributes. A reference HUB including one or more schemas can be used to analyze data flows, and further classify or make recommendations such as, for example, transformations enrichments, filtering, or cross-entity data fusion of an input data.

[00071] In accordance with an embodiment, the system provides a programmatic interface, referred to herein in some embodiments as a foreign function interface, by which a user or third-party can define a service, functional and business types, semantic actions, and patterns or predefined complex data flows based on functional and business types, in a declarative manner, to extend the functionality of the system.

[00072] In accordance with an embodiment, the system can provide data governance functionality such as, for example, provenance (where a particular data came from), lineage (how the data was acquired/processed), security (who was responsible for the data), classification (what is the data about), impact (how impactful is the data to a business), retention (how long should the data live), and validity (whether the data should be excluded/included for analysis/processing), for each slice of data pertinent to a particular snapshot in time; which can then be used in making lifecycle decisions and dataflow recommendations.

[00073] In accordance with an embodiment, the system can be implemented as a service, for example as a cloud service provided within a cloud-based computing environment; and can serve as a single point of control for the design, simulation, deployment, development, operation, and analysis of data for use with software applications; including enabling data input from one or more sources of data (for example, in an embodiment, an input HUB); providing a graphical user interface that enables a user to specify an application for the data; and, scaling the data depending on an intended destination, use, or target for the data (for example, in an embodiment, an output HUB).

[00074] In accordance with an embodiment, as used herein, the terms “input” and “output” when used with reference to a particular HUB are provided merely as labels to reflect the apparent flow of data in particular use cases or examples, and are not intended to be

restrictive as to the type or function of the particular HUB.

[00075] For example, in accordance with an embodiment, an input HUB that operates as a source of data can also, at the same or at another time, operate as an output HUB or target to receive the same, or another data, and vice versa.

5 **[00076]** Additionally, while for purposes of illustration, several of the examples described herein illustrate the use of an input HUB and an output HUB; in accordance with an embodiment, in practical implementations, a data integration or other computing environment can include a plurality of such HUBs, at least some of which are acting both as input HUBs and/or output HUBs.

10 **[00077]** In accordance with an embodiment, the system enables rapid development of software applications in large, e.g., cloud-based, computing environments, where data models may evolve rapidly, and where features such as, for example, search, recommendations, or suggestions are valuable business requirements. In such environments, the combination of artificial intelligence (AI) and semantic search empowers users to accomplish more with their
15 existing systems. For example, integration interactions, such as attribute-level mappings, can be recommended based on an understanding of metadata, data, and a user's interactions with the system.

[00078] In accordance with an embodiment, the system can also be used to suggest complex cases, for example interesting dimensional edges, which can be used to analyze
20 information, and which enable users to discover hitherto unknown facts within their data.

[00079] In some embodiments, the system provides a graphical user interface that enables the automation of manual tasks (e.g., recommendations or suggestions), and leverages machine learning and probabilistic knowledge federation, to provide a useful context for users, and allow for discovery and semantics-driven solutions, for example, the creation of
25 a data warehouse, scaling of services, data preparation and enrichment, and design and monitoring of software applications.

[00080] In accordance with various embodiments, the system can include or utilize some or all of the following features:

[00081] Design-Time System: In accordance with an embodiment, a computational
30 environment that enables the design, creation, monitoring, and management of software applications (for example, a dataflow application, pipeline, or Lambda application), including the use of an, e.g., data AI subsystem, that provides machine learning capabilities.

[00082] Run-Time System: In accordance with an embodiment, a computational
35 environment that enables the execution of software applications (for example, a dataflow application, pipeline, or Lambda application), and that receives input from, and provides recommendations to, a design-time system.

[00083] Pipeline: In accordance with an embodiment, a declarative means of defining a processing pipeline, having a plurality of stages or semantic actions, each of which corresponds to a function such as, for example, one or more of filtering, joining, enriching, transforming, or fusion of an input data, for preparation as an output data. A dataflow software application, or dataflow application, representing a data flow in, e.g., DFML. In accordance with an embodiment, the system supports a declarative pipeline design that can use a same code base (e.g., with a Spark runtime platform) for both batch (historical) and real time (streaming) data processing; and also supports the building of pipelines or applications that can operate on real time data streams, for real time data analytics. Data reprocessing due to pipeline design changes can be handled through rolling upgrades of deployed pipelines. In accordance with an embodiment, a pipeline can be provided as a Lambda application that can accommodate the processing of real time data and batch data within distinct batch and real time layers.

[00084] HUB: In accordance with an embodiment, a data source or target (cloud or on-premise) comprising datasets or entities. A data source, that can be introspected, from which data can be consumed or published to, and which comprises datasets or entities, that have attributes, semantics or relationships with other datasets or entities. Examples of HUBs include streaming data, telemetric, batch-based, structured or unstructured, or other types of data sources. Data can be received from one HUB, associated with a source dataset or entity, and mapped to a target dataset or entity at the same or another HUB.

[00085] System HUB: In accordance with an embodiment, a system HUB can operate as a knowledge source to store profile information and other metadata that can be associated with other HUBs, and datasets or entities, in those HUBs, and can also operate in the manner of a regular HUB as a source or recipient of data to be processed. A central repository in, e.g., DFML where metadata and state of the system is managed.

[00086] Dataset (Entity): In accordance with an embodiment, a data structure comprising attributes (e.g., columns), that can be owned by or otherwise associated with one or more HUBs, for example a database table, view, file or API. In accordance with an embodiment, one or more business entities, for example customer records, that can function as a semantic business type, and are stored as data components, such as, for example, tables, within a HUB. Datasets or entities can have relationships to other datasets or entities; together with attributes, such as, for example, a column in a table; and a data type, such as, for example, string or integer. In accordance with an embodiment, the system supports a schema agnostic processing of all types of data (including for example, structured, semi-structured, or unstructured data) during, for example, enrichment, preparation, transformation, model training, or scoring operations.

[00087] Data AI Subsystem: In accordance with an embodiment, a component of a system, such as, for example, a Data AI system, responsible for machine-learning and semantic related functions, including one or more of search, profiling, providing a recommendation engine, or support for auto-mapping. The data AI subsystem can support, through an event coordinator, the operations of design-time systems, for example a software development component, e.g., Lambda Studio, and can provide recommendations based on the continued processing of data by dataflow applications (e.g., pipelines, Lambda applications), for example to recommend the modification of an existing, e.g., pipeline, to take better advantage of the data being processed. A data AI subsystem can analyze amounts of input data, and continuously update a domain-knowledge model. During the processing of a dataflow application (e.g., pipeline), each stage of the, e.g., pipeline, can proceed, based on recommended alternatives or options provided by the data AI subsystem, the updated domain model, and inputs from users, e.g., to accept or reject a recommended semantic action.

[00088] Event Coordinator: In accordance with an embodiment, an event-driven architecture (EDA) component that operates between the design-time and run-time systems, to coordinate events related to the design, creation, monitoring, and management of dataflow applications (e.g., pipelines, Lambda applications). For example, the event coordinator can receive a published notification of data from a HUB (e.g., new data conforming to a known data type), normalize the data from that HUB, and provide the normalized data to a set of subscribers, for use by, e.g., pipelines, or other downstream consumers. The event coordinator can also receive notification of state transactions within the system, for use in logging, or lineage tracking, including the creation of temporal slices; and schema evolution.

[00089] Profiling: In accordance with an embodiment, the operation of extracting a sample of data from a HUB, in order to profile the data provided by that HUB, and the datasets or entities, and attributes within that HUB; together with determining metrics associated with sampling the HUB, and updating the metadata associated with a HUB to reflect a profile of the data in that HUB.

[00090] Software Development Component (Lambda Studio): In accordance with an embodiment, a design-time system tool that provides a graphical user interface to enable a user to create, monitor, and manage the lifecycle of a pipeline or Lambda application as a pipeline of semantic actions. A graphical user interface (UI, GUI) or studio that allows users to design, e.g., pipelines, Lambda applications.

[00091] Semantic Action: In accordance with an embodiment, a data transformation function, for example a relational algebraic operation. An action that can be performed by dataflow applications (e.g., pipelines, Lambda applications) on a dataset or entity within a HUB, for projection onto another entity. A semantic action operates as a higher-order function that

can be used across different models, or HUBs, that can receive a dataset input, and produce a dataset output. Semantic actions can include mappings, and can be continuously updated by, e.g., the data AI subsystem, in response to processing data as part of an e.g., pipeline, or Lambda application.

5 **[00092]** Mapping: In accordance with an embodiment, a recommended mapping of a semantic action between a first (e.g., source) dataset or entity, and another (e.g., target) dataset or entity, provided by, e.g., the data AI subsystem, and made accessible through the design-time system, e.g., via the software development component, Lambda Studio. For example, a data AI subsystem can provide auto-mapping as a service, wherein the auto-

10 mapping can be driven by a metadata, schema, and statistical profiling of a dataset, based on a machine learning analysis of a metadata associated with a HUB or data input, and a profiling of the data itself.

[00093] Pattern: In accordance with an embodiment, a pattern of semantic actions that can be performed by a dataflow application (e.g., a pipeline, Lambda application). A template

15 can be used to provide a definition of a pattern that can be reused by other applications. A data flow representing a logical flow of data and associated transformations usually associated with business semantics and processes.

[00094] Policies: In accordance with an embodiment, a set of policies that control how dataflow applications (e.g., pipelines, Lambda applications) are scheduled, which users or

20 components can access which HUBs and semantic actions, and how data should be aged, or other considerations. A configuration setting that defines how an, e.g., pipeline, is, for example, scheduled, executed, or accessed.

[00095] Application Design Services: In accordance with an embodiment, provides data flow, e.g., pipeline, Lambda application specific services such as, for example, validation,

25 compiling, packaging, deployment, to other, e.g., DFML services (e.g., UI, system façade). A design-time system component that validates the pipeline of an, e.g., pipeline, or Lambda application in the software development component, e.g., Lambda Studio (e.g., its inputs and outputs), persists the pipeline, and controls the deployment of the pipeline, Lambda application, to the system (e.g., to a Spark cluster) for execution, and thereafter can be used

30 in managing the lifecycle or state of the application.

[00096] Edge Layer: In accordance with an embodiment, a layer that collects and forwards data to the scalable I/O layer, e.g., as a store and forward layer. A run-time system component that includes one or more nodes that can receive data, e.g., via a gateway that is

35 accessible to the Internet, and that includes security and other features that support secured access to the, e.g., Data AI system.

[00097] Compute Layer: In accordance with an embodiment, an application execution

and data processing layer (e.g., Spark). A run-time system component operating as a distributed processing component, for example, a Spark cloud service, cluster of compute nodes, collection of virtual machines, or other components or nodes, for use in the execution of, e.g., pipelines, Lambda applications. In a multitenant environment, nodes within the compute layer can be allocated to tenants for use in executing pipelines or Lambda applications by those tenants.

[00098] Scalable Input/Output (I/O) Layer: In accordance with an embodiment, provides scalable data persistence and access layer structured as topics and partitions (e.g., Kafka). A run-time system component providing a queue or other logical storage that allows data to be moved within the system, and shared between the various components of the system, for example, a Kafka environment. In a multitenant environment, the scalable I/O layer can be shared among multiple tenants.

[00099] Data Lake: In accordance with an embodiment, a repository for the persistence of information from the system HUB, or other components. A repository of data in, e.g., DFML, usually normalized or processed by, e.g., pipelines, Lambda applications, and consumed by other pipelines, Lambda applications or a publish layer.

[000100] Registries: In accordance with an embodiment, one or more repositories of information, e.g., for storage of functional and business types, that are used in the decomposition of, e.g., pipelines, Lambda applications into their functional components.

[000101] DataFlow Machine Learning (DFML): In accordance with an embodiment, a data integration, data flow management system that leverages machine learning (ML) to assist in the building of complex dataflow applications.

[000102] Metadata: In accordance with an embodiment, an underlying definition, description of the datasets or entities, and attributes and their relationships. It can also be descriptive data about an artifact in, e.g., DFML.

[000103] Data: In accordance with an embodiment, an application data represented by the datasets or entities. They can be batch or stream. For example, customers, orders, or products.

[000104] System Façade: In accordance with an embodiment, a unified API Layer to access the functional capabilities of an e.g., DFML event-driven architecture.

[000105] Data AI subsystem: In accordance with an embodiment, provides artificial intelligence (AI) services including but not limited to, for example, search, auto-map, recommendation, or profiling.

[000106] Streaming Entity: In accordance with an embodiment, continuous input of data and near real time processing and output requirements, which may support an emphasis on velocity of data.

[000107] Batch Entity: In accordance with an embodiment, scheduled or on-request ingestion of data, which can be characterized by an emphasis on volume.

[000108] Data Slice: In accordance with an embodiment, a partition of data usually marked by time.

5 **[000109]** Rule: In accordance with an embodiment, represents a directive governing artifacts in, e.g., DFML, for example data rules, relationship rules, metadata rules and complex or hybrid rules.

[000110] Recommendation (Data AI): In accordance with an embodiment, a proposed course of action usually represented by one or more semantic actions or fine grained directives
10 to assist in the design of, e.g., pipelines, Lambda applications.

[000111] Search (Data AI): In accordance with an embodiment, a semantic search in, e.g., DFML, characterized by the context and the intent of the user to return relevant artifacts.

[000112] Auto-map (Data AI): In accordance with an embodiment, a type of recommendation that shortlists the candidate source or target datasets or entities, to be used
15 in a data flow.

[000113] Data Profiling (Data AI): In accordance with an embodiment, a collection of several metrics, for example, min value, max value, inter-quartile range, or sparsity, that characterize data in an attribute belonging to a dataset or entity.

[000114] Action Parameter: In accordance with an embodiment, a reference to datasets on which the semantic actions are performed. For example, the parameters to an equi-join in
20 an, e.g., pipeline, Lambda application.

[000115] Foreign Function Interface: In accordance with an embodiment, a mechanism to register and invoke services (and semantic actions) as part of the, e.g., DFML Lambda application Framework. It can be used to extend the capabilities or transformation vocabulary
25 in, e.g., DFML.

[000116] Service: In accordance with an embodiment, an owning artifact in, e.g., DFML, of a collection of semantic actions that can be characterized by a data integration stage (for example, preparation, discovery, transformation, or visualization).

[000117] Service Registry: In accordance with an embodiment, a repository of services, their semantic actions and other instance information.
30

[000118] Data Lifecycle: In accordance with an embodiment, the stages in the use of data within, e.g., DFML, starting from ingestion and ending at publication.

[000119] Metadata Harvesting: In accordance with an embodiment, the collection of metadata and sample data for profiling, usually after the registration of a HUB.

35 **[000120]** Normalize Pipelines: In accordance with an embodiment, the standardization of data in a certain format to facilitate consumption by, e.g., pipelines, Lambda applications.

[000121] Monitoring: In accordance with an embodiment, identifying, measuring and evaluating the performance of an, e.g., pipeline, Lambda application.

[000122] Ingest: In accordance with an embodiment, the intake of data through the edge layer in, e.g., DFML.

5 [000123] Publish: In accordance with an embodiment, the writing data to a target end point from, e.g., DFML.

Data AI System

[000124] Figure 1 illustrates a system for providing data flow artificial intelligence, in
10 accordance with an embodiment.

[000125] As illustrated in Figure 1, in accordance with an embodiment, a system, e.g., data AI system 150, can provide one or more services for processing and transforming data such as, for example, business data, consumer data, and enterprise data, including the use of machine learning processing, for use with a variety of computational assets such as, for
15 example, databases, cloud data warehouses, storage systems, or storage services.

[000126] In accordance with an embodiment, computational assets can be cloud-based, enterprise-based, or on-premise or agent-based. Various elements of the system can be connected by one or more networks 130.

[000127] In accordance with an embodiment, the system can include one or more input
20 HUBs 110 (e.g., sources of data, data sources), and output HUBs 180 (e.g., targets of data, data targets).

[000128] In accordance with an embodiment, each input HUB, e.g., HUB 111, can include a plurality of (source) datasets or entities 192.

[000129] In accordance with an embodiment, examples of input HUBs can include a
25 database management system (DB, DBMS) 112 (e.g., an on-line transaction processing system (OLTP), business intelligence system, or an on-line analytical processing system (OLAP)). In such examples, the data provided by a source such as, for example, a database management system can be structured or semi-structured.

[000130] In accordance with an embodiment, other examples of input HUBs can include
30 a cloud store/object store 114 (e.g., AWS S3, or another object store), which can be a click-stream source or an object bucket with unstructured data; data clouds 116 (e.g., a third-party cloud); a streaming data source 118 (e.g., AWS Kinesis, or another streaming data source), or other input sources 119.

[000131] In accordance with an embodiment, input HUBs can include a data source into
35 which data is received from, e.g., an Oracle Big Data Prep (BDP) service.

[000132] In accordance with an embodiment, the system can include one or more output

HUBs 180 (e.g., output destinations). Each output HUB, e.g., HUB 181, can include a plurality of (target) datasets or entities 194.

[000133] In accordance with an embodiment, examples of output HUBs can include a public cloud 182, a data cloud 184 (e.g., AWS and Azure), an on-premise cloud 186, or other
5 output targets 187. Data output provided by the system can be produced for dataflow applications (e.g., pipelines, Lambda applications) accessible at an output HUB.

[000134] In accordance with an embodiment, examples of public cloud can be, for example, an Oracle Public Cloud, including services such as, for example, Big Data Prep Cloud Service, Exadata Cloud Service, Big Data Discovery Cloud Service, and Business Intelligence
10 Cloud Service.

[000135] In accordance with an embodiment, the system can be implemented as a unified platform for streaming and on-demand (batch) data processing delivered as a service (e.g., as a software-as-a-service) to users, providing scalable, multi-tenant data processing for multiple input HUBs. Data is analyzed in real-time, using machine learning techniques and
15 visual insights and monitoring provided by a graphical user interface as part of the service. Data sets can be fused from multiple input HUBs for output to an output HUB. For example, through data processing services provided by the system, data can be generated for a data warehouse and population in one or more output HUBs.

[000136] In accordance with an embodiment, the system provides declarative and
20 programming topologies for transformation, enrichment, routing, classification, and blending of data; and can include a design-time system 160 and a run-time system 170. Users can create applications, such as, for example, dataflow applications (e.g., pipelines, Lambda applications) 190, designed to perform data processing.

[000137] In accordance with an embodiment, the design-time system can enable users
25 to design dataflow applications, define data flows, and define data for data flow processing. For example, design-time system can provide a software development component 162 (referred to herein in an embodiment as Lambda Studio) that provides a graphical user interface for creation of dataflow applications.

[000138] For example, in accordance with an embodiment, using the software
30 development component, a user can specify input HUBs and output HUBs for creating a data flow for an application. The graphical user interface can present interfaces for services for data integration, which enables a user to create, manipulate, and manage data flow for an application, including the ability to monitor and manage a data flow pipeline dynamically, such as, for example, viewing data lineage and performing forensic analysis.

[000139] In accordance with an embodiment, the design-time system can also include
35 application design services 164 for deploying dataflow applications into run-time system.

[000140] In accordance with an embodiment, the design-time system can also include one or more system HUBs 166 (e.g., a metadata repository) for storing metadata for processing data flows. The one or more system HUBs can store samples of data, such as, for example, data types including functional and business data types. The information in the system HUB can be used to perform one or more of the techniques disclosed herein. A data lake 167 component, can operate as a repository for the persistence of information from the system HUB.

[000141] In accordance with an embodiment, the design-time system can also include a data artificial intelligence (AI) subsystem 168 to perform operations for data artificial intelligence processing. The operations can include the use of ML techniques, for example search and retrieval. The data AI subsystem can sample data for generating metadata for the system HUB.

[000142] In accordance with an embodiment, for each input HUB, the data AI subsystem can perform schema object analysis, metadata analysis, sample data, correlation analysis, and classification analysis. The data AI subsystem can provide rich data to dataflow applications by running continuously on data that is input; and can provide recommendations, insights, and type inductions to, for example, pipelines, Lambda applications.

[000143] In accordance with an embodiment, the design-time system enables a user to create policies, artifacts and flows that define the functional needs of the use case.

[000144] For example, in accordance with an embodiment, the design-time system can provide a graphical user interface to create a HUB to ingest data and define the ingest policy which can be time based or as needed by the related data flows. Upon selecting an input HUB, data can be sampled from the input HUB to profile the sources, such as, for example, performing metadata query, obtaining samples, and obtaining user defined input. The profile can be stored in system HUB. The graphical user interface enables multiple sources to be joined for defining a dataflow pipeline. This can be done by creating a script or by using the guided editor by which data can be visualized at each step. The graphical user interface can provide access to a recommendation service that suggests how the data could be, e.g., corrected, enriched, or joined.

[000145] In accordance with an embodiment, during design time, application design services can suggest a suitable structure to analyze the resulting content. The application design services can use a knowledge service (functional type classification) to suggest measures and related dimension hierarchies. Once that is completed, the design-time system can recommend the data flows needed to take the blended data from earlier pipelines and populate a dimensional target structure. Based on the dependency analysis it can also derive and generate the orchestration flow to load / refresh the target schema. For forward-

engineering use cases, the design-time system can also generate a HUB to host a target structure and create the target schema.

[000146] In accordance with an embodiment, the run-time system can perform processing during run-time of a service for data processing.

5 **[000147]** In accordance with an embodiment, in the run-time, or operation mode, the policy and flow definitions created by the user are applied and/or executed. For example, such processing can include invoking ingest, transform, model and publish services, to process data in the pipeline.

[000148] In accordance with an embodiment, the run-time system can include an edge
10 layer 172, a scalable input/output (I/O) layer 174, and a distributed processing system or compute layer 176. During run-time (e.g., when data is ingested from one or more input HUBs 110), data can be received by edge layer for events that cause the data to be generated.

[000149] In accordance with an embodiment, an event coordinator 165 operates
15 between the design-time and run-time systems, to coordinate events related to the design, creation, monitoring, and management of the dataflow applications (e.g., pipelines, Lambda applications).

[000150] In accordance with an embodiment, an edge layer sends the data to the scalable I/O layer for routing the data to distributed processing system or compute layer.

[000151] In accordance with an embodiment, a distributed processing system or
20 compute layer can implement a pipeline process (on a per-tenant basis) to process data for output. The distributed processing system can be implemented using, for example, Apache Spark and Alluxio; and can sample data into a data lake before outputting the data to an output HUB. The distributed processing system can communicate with the scalable I/O layer for activation of data for processing.

25 **[000152]** In accordance with an embodiment, the data AI system, including some or all of the above components, can be provided at, or executed by, one or more computers including, for example, one or more processors (CPU), memory, and persistent storage devices (198).

30 **Event-Driven Architecture**

[000153] As described previously, in accordance with an embodiment, the system can include an event-driven architecture (EDA) component or event coordinator that operates between the design-time and run-time systems, to coordinate events related to the design, creation, monitoring, and management of dataflow applications (e.g., pipelines, Lambda
35 applications).

[000154] **Figure 2** illustrates an event-driven architecture including an event coordinator

for use with a system, in accordance with an embodiment.

[000155] As illustrated in Figure 2, in accordance with an embodiment, the event coordinator can include an event queue 202 (e.g., Kafka), event bootstrapper service 204 (e.g., ExecutorService), and event configuration publisher / event consumer 206 (e.g., DBCS).

5 **[000156]** In accordance with an embodiment, events received at a system façade 208 (e.g., an Events API Extension) can be communicated by one or more event broker(s) 210, e.g., Kafka consumer(s), to various components of the system. For example, data and/or events such as, for example, external data 212 (e.g., S3, OSCS, or OGG data), or inputs from a graphical user interface 214 (e.g., a browser, or DFML UI) can be communicated via the event coordinator, to other components such as, for example, an application runtime 216, data lake, system HUB, data AI subsystem, application design services, as described previously; and/or to ingest 220, publish 230, scheduling 240, or other components.

10 **[000157]** In accordance with an embodiment, event brokers can be configured as consumers of stream events. The event bootstrapper can start a number of configured event brokers to process events on behalf of registered subscribers; wherein each event broker delegates the processing of the event to the registered callback endpoint for processing a given event. The event coordinator enables registering event types; registering eventing entities; registering events; and registering subscribers. Table 1 provides an example of various event objects, including a publish event, and a subscribe event.

20

Publish Event	A 2-tuple (event Type, eventing Entity) object that registers the publication of an event by an eventing entity.
Subscribe Event	A 2-tuple (event, eventing Entity) object that registers the subscription to an event published by another eventing entity (publisher).

Table 1

25 **Event Types**

[000158] In accordance with an embodiment, event types define a state change of an event that has significance for the system, such as, for example, the creation of a HUB, modification of a dataflow application, e.g., pipeline, Lambda application, ingesting data for a dataset or entity, or publishing data to a target HUB. An example data format, and examples
 30 of various event types are described below and in Table 2:

```

{
  "Id": "" , // generated at the time of creation
  "Name": "Hub - Creation"
  "Type": "User Event"
5 }
    
```

POST /eventTypes	Creating a new eventType.
PUT /eventType/{eventid}	Modify an eventType with the given Id.
GET /eventTypes	Retrieve all eventTypes in the system.
GET /eventTypes/{eventTypeId}	Retrieve the representation of the eventType with the given Id.
GET /eventTypes/{eventid}/publishers	Retrieve all the publishers for this eventType. There can be more than one eventing entity registered to publish events of this type, in which case all distinct eventing entities publishing event of this type will be returned.
GET /eventTypes/{eventid}/subscribers	Retrieve all the subscribers for this eventType. There can be more than one eventing entity registered as subscribers to events of this type, in which case all distinct eventing entities subscribing to the event of this type will be returned.
DELETE /eventTypes/{eventid}	Delete an event with a given Id.

Table 2

10 **Eventing Entities**

[000159] In accordance with an embodiment, an eventing entity can be a publisher and/or subscriber of events. For example, an eventing entity can register to publish one or more events and/or be consumers of one or more events, including registering an endpoint or callback URL that will be used for notifying or sending acknowledgement on publish and delegating processing for subscribed events. Examples of eventing entities can include the metadata service, ingest service, system HUB artifacts, and pipelines, Lambda applications. An example data format, and examples of various eventing entities are described below and in Table 3:

```

{
  "Id":"" , // generated at the time of creation
  "Name": "Metadata Service",
  "endPointURL":"localhost:9010/processEvent",
5  "entityType":"DFMLService"
}
    
```

POST /eventingEntities	Creates a new eventing Entity.
PUT /eventingEntities/{entityId}	Modify an existing eventing Entity identified by Id.
GET /eventingEntities	Retrieve all registered eventing Entities.
GET /eventingEntities/{entityId}	Retrieve the representation of the eventing entity with the given Id.
GET /eventingEntities/{entityId}/eventsPublished	Retrieve all events registered for publication by this eventing Entity.
GET /eventingEntities/{entityId}/eventsSubscribed	Retrieve all events registered for subscription by this eventing Entity.
DELETE /eventingEntities/{entityId}	Delete this evening entity from the system.

Table 3

10

Events

[000160] In accordance with an embodiment, an event is an instance of an event type that is associated with an eventing entity registered as a publisher; and can have subscribers (eventing entities). For example, a metadata service can register a HUB creating event for publication; and can publish one or more of event instances (once each for HUBs created) for this event. Examples of various events are described in Table 4:

15

POST /events	Creates a new event to be published by an eventing Entity. { "Id":"" ,
--------------	--

	<pre>"acknowledgeURL": "localhost:9010/eventAcknowledge", "onProcessingOf":"/eventType/{eventId} ", "eventType":"/eventType/{eventId}", "eventingEntity":"/eventingEntity/{entitld} " }</pre>
PUT /events/{eventid}	Modifies an existing event registered for publication by an eventing entity.
POST /events/{eventid}/publish	<p>Publishes an event instance enqueueing the event for consumption.</p> <pre>{ "event_type": "data", "subtype": "publication", "state": "ready", "context": {"eventContextId":"{eventId}","accessTo ken":""} "message": { _actual event data goes here_ } }</pre>
GET /events	Retrieve all events registered in the system.
GET /events/{eventide}	Retrieve the event with the given Id.
GET /events/{eventide}/publisher	Retrieve the publisher of this event.
GET /events/{eventide}/subscribers	Retrieve the subscribers for this event.
POST /events/{eventid}/subscribers	<p>Registering a subscriber for an event.</p> <pre>{ "Id":""," // generated at the time of creation "processingURL": "localhost:9010/eventProcess",</pre>

	<pre>"subscribingEvent":"/events/{eventId}", "callbackMethod":"REST", "subscriberentity":"/eventingEntity/{entitl d}" }</pre>
PUT /events/{eventId}/subscribers/{subscriberId}	Modify the properties of the subscriber for this event.
DELETE /events/{eventId}/subscribers/{subscriberId}	Delete the subscriber for this event.
DELETE /events/{eventid}	Delete the event.

Table 4

Examples

5 **[000161]** In accordance with an embodiment, the following examples illustrate creating an event type; registering a publish event; registering subscribers; publishing an event; getting event types; getting publishers for an event type; and getting subscribers for an event type.

POST http://den00tnk:9021/dfml/service/eventType

10

```
{ "name": "DATA_INGESTED",
  "type": "SystemEvent"
}
```

15 **[000162]** Which returns a universally unique ID (UUID), for example “8e87039b-a8b7-4512-862c-fdb05b9b8888”. Eventing objects can publish or subscribe to events in the system. Service endpoints such as, for example, ingest service, metadata service, and application design service can be publish and subscribe events with static endpoints for acknowledge, notification, error, or processing. DFML artifacts (e.g., DFMLEntity, DFMLLambdaApp, 20 DFMLHub) can also be registered as eventing objects; instances of these types can publish and subscribe to events without registering as eventing objects.

POST http://den00tnk:9021/dfml/service/eventEntity

25

```
{
  "name": "DFMLEntity",
```

```

    "endPointURL": "localhost:9010/<publisherURL>",

    "notificationEndPointURL": "http://den00tnk:9021/<publisherURL>/notification",
    "exceptionEndPointURL": "http://den00tnk:9021/<publisherURL>/exception",
    "errorEndPointURL": "http://den00tnk:9021/<publisherURL>/error",
    "entityType": "DFMLEntity"
  }

```

10 **[000163]** The following example registers the DFMLLambdaApps (type) as an eventing Object.

```

{
  "name": "DFMLLambdaApps",
  "endPointURL": "localhost:9010/<publisherURL>",

  "notificationEndPointURL": "http://den00tnk:9021/<publisherURL>/notification",
  "exceptionEndPointURL": "http://den00tnk:9021/<publisherURL>/exception",
  "errorEndPointURL": "http://den00tnk:9021/<publisherURL>/error",
  "entityType": "DFMLLambdaApps"
}

```

25 **[000164]** For eventing entities that are of type HUB, Entity and LambdaApp, the <publisherURL> can be annotated in REST endpoint URL's, and the event-driven architecture will derive the actual URL by replacing the DFML artifact instance URL. For example, if a notificationEndpointURL is registered as http://den00tnk:9021/<publisherURL>/notification and the publisher URL specified as part of the message is hubs/1234/entities/3456, then the URL invoked for notification will be http://den00tnk:9021/ hubs/1234/entities/3456 /notification.

30 The POST returns a UUID; for example "185cb819-7599-475b-99a7-65e0bd2ab947".

Registering a Publish Event

[000165] In accordance with an embodiment, a publish event can be registered as:

```

35 POST http://den00tnk:9021/dfml/service/event

{
  "acknowledgeURL": "http://den00tnk:9021/<publisherURL>/acknowledge",
  "onProcessingOf": "/eventType/{eventId}",

```

```
"eventType": "7ea92c868e87039b-a8b7-4512-862c-fdb05b9b8888",
"publishingEntity": "185cb819-7599-475b-99a7-65e0bd2ab947"
}
```

5 **[000166]** The eventType refers to the UUID returned for registering the event type DATA_INGESTED and the publishingEntity refers to the DFMLEntity type registered as eventing Object. The registration returns a UUID, for example "2c7a4b6f-73ba-4247-a07a-806ef659def5".

10 Registering Subscribers

[000167] In accordance with an embodiment, a subscriber can be registered as:

```
POST http://den00tnk:9021/dfml/service/event/2c7a4b6f-73ba-4247-a07a-
806ef659def5/subscribers
```

15

[000168] The UUID returned from the publish event registration is used in the path segment for registering a subscriber.

```
{
20 "processingURL": "http://den00tnk:9021/dfml/service/eventType/process3",
"callbackMethod": "SYNC_POST",
"subscriberEntity": "7599916b-baab-409c-bfe0-5334f111ef41",
"publisherURL": "/hubs/1234/entities/3456",
"publishingObjectType": "DFMLEntity",
25 "subscribingObjectType": "DFMLLambdaApps",
"subscriberURL": "/lambdaApps/123456"
}
```

30 **[000169]** The publisherURL and publishingObjectType refers to the instance and type of the publisher object. Here the dataflow (e.g, Lambda) application, in specifying a URI /lambdaApps/123456 is interested in subscribing to the DATA_INGESTED event from the entity /hubs/1234/entities/3456. The registration returns a UUID, for example "1d542da1-e18e-4590-82c0-7fe1c55c5bc8".

35 Publishing an Event

[000170] In accordance with an embodiment, an event can be published as:

POST http://den00tnk:9021/dfml/service/event/publish

```

{
  "event_type": "DATA_AVAILABLE",
  "subtype": "publication",
  5  "state": "ready",
  "eventType": "2c7a4b6f-73ba-4247-a07a-806ef659def5",
  "publisherURL": "dfml/service/eventType",
  "message": { "id": "1234",
    "descr": "something happened here testing this and this again"
  }
10 }
}

```

[000171] The publisherURL is used if the publishing object is one of DFMLEntity, DFMLHub or DFMLLambdaApps, and is used to check the instance of the eventing object publishing a message for which the subscriber is enlisted. The publisher URL is also used to derive the notification URL when a subscriber successfully processes the message. The publication returns a message body that was part of the published event.

Getting Event Types

20 **[000172]** In accordance with an embodiment, event types can be determined as:

GET http://den00tnk:9021/dfml/service/eventType

```

{
25  "eventTypes": [
    {
      "Id": "8e87039b-a8b7-4512-862c-fdb05b9b8888",
      "name": "DATA_INGESTED",
      "type": "SystemEvent",
      30  "createdBy": " ",
      "updatedBy": " ",
      "description": " ",
      "typeQualifier": " ",
      "resourceType": " ",
      35  "verb": " ",
      "operationType": " ",
      "status": " ",
      "annotation": " "
    },
  ],

```

```

    {
      "Id": "7ea92c86-8db5-42d6-992a-2578a6d025ce",
      "name": "DATA_AVAILABLE",
      "type": "SystemEvent",
5     "createdBy": " ",
      "updatedBy": " ",
      "description": " ",
      "typeQualifier": " ",
      "resourceType": " ",
10    "verb": " ",
      "operationType": " ",
      "status": " ",
      "annotation": " "
    }
15  ]
  }

```

[000173] Getting Publishers for an Event Type

20 GET http://den00tnk:9021/dfml/service/eventType/7ea92c86-8db5-42d6-992a-2578a6d025ce/publishers

```

{
  "eventingObjects": [
25   {
      "Id": "185cb819-7599-475b-99a7-65e0bd2ab947",
      "name": "DFMLEntity",
      "entityType": "DFMLEntity",
      "endpointURL": "localhost:9010/<publisherURL>",
30    "notificationEndpointURL":
"http://den00tnk:9021/<publisherURL>/notification",
      "exceptionEndpointURL":
"http://den00tnk:9021/<publisherURL>/exception",
      "errorEndpointURL": "http://den00tnk:9021/<publisherURL>/error",
35    "acknowledgeEndpointURL": " ",
      "description": " ",
      "entityQualifier": " ",
      "status": " ",
      "annotation": " "
40   }

```

```

]
}

```

Getting Subscribers for an Event Type

5 **[000174]** In accordance with an embodiment, subscribers for an event type can be determined as:

```

GET http://den00tnk:9021/dfml/service/eventType/7ea92c86-8db5-42d6-992a-
2578a6d025ce/subscribers

```

10

```

{
  "eventingObjects": [
    {
      "Id": "7599916b-baab-409c-bfe0-5334f111ef41",
      "name": "DFMLLambdaApps",
      "entityType": "DFMLLambdaApps",
      "endPointURL": "localhost:9010/<publisherURL>",
      "notificationEndpointURL":
      "http://den00tnk:9021/<publisherURL>/notification",
      "exceptionEndpointURL":
      "http://den00tnk:9021/<publisherURL>/exception",
      "errorEndpointURL": "http://den00tnk:9021/<publisherURL>/error",
      "acknowledgeEndpointURL": " ",
      "description": " ",
      "entityQualifier": " ",
      "status": " ",
      "annotation": " "
    }
  ]
}

```

30

[000175] The illustrations described above are provided by way of example, to illustrate a particular embodiment of an event coordinator, event types, eventing entities, and events. In accordance with other embodiments, other types of EDAs can be used to provide communication within the system operating between the design-time and run-time systems, to coordinate events related to the design, creation, monitoring, and management of dataflow applications, and other types of event types, eventing entities, and events can be supported.

35

DataFlow Machine Learning (DFML) Flow

[000176] As described previously, in accordance with various embodiments, the system can be used with a data integration or other computing environment, that leverages machine learning (ML, DataFlow Machine Learning, DFML), for use in managing a flow of data (dataflow, DF), and building complex dataflow software applications (e.g., dataflow applications, pipelines, Lambda applications).

[000177] **Figure 3** illustrates the steps in a data flow, in accordance with an embodiment.

[000178] As illustrated in **Figure 3**, in accordance with an embodiment, the processing of a DFML data flow 260 can include a plurality of steps, including an ingest step 262, during which data is ingested from various sources, for example, Salesforce (SFDC), S3, or DBaaS.

[000179] During a data preparation step 264, the ingested data can be prepared, for example by being de-duplicated, standardized, or enriched.

[000180] During a transform step 266, the system can perform one or more merges, filters, or lookups at datasets, to transform the data.

[000181] During a model step 268, one or more models are generated, together with mappings to the models.

[000182] During a publish step 270, the system can publish models, specify policies and schedules, and populate target data structures.

[000183] In accordance with an embodiment, the system supports the use of a search/recommend functionality 272 throughout each of its data preparation, transform, and model steps. A user can interact with the system through a set of well-defined services that encapsulates the breadth of functional capabilities in data integration frameworks. This set of services define a logical view of the system. For example, in a design mode, the user can create policies, artifacts, and flows that define the functional needs of a particular use case.

[000184] **Figure 4** illustrates an example of a data flow that includes multiple sources, in accordance with an embodiment.

[000185] As illustrated in the example data flow 280 shown in **Figure 4**, in accordance with an embodiment, the requirement is to take content from a plurality of sources 282, here indicated as SFDC, and FACS (Fusion Apps Cloud Service), together with some files in OSCS (Oracle Storage Cloud Service); blend that information together in a manner that can be used to analyze the desired content; derive a target cube and dimensions; map the blended content to the target structures; and make this content available along with the dimensional model to an Oracle Business Intelligence Cloud Service (BICS) environment; including the use of ingest, transform 266A/266B, model, orchestrate 292, and deploy 294 steps.

[000186] The illustrated example is provided to illustrate the techniques described herein; the functionality described herein is not limited to use with these particular data sources.

[000187] In accordance with an embodiment, during the ingest step, in order to access and ingest the SFDC content, a HUB is created in the data lake to receive this content. This can be performed, for example, by selecting an SFDC adapter for the relevant access mode (JDBC, REST, SOAP), creating the HUB, providing a name, and defining an ingest policy
5 which could be time based or as needed by the related data flows.

[000188] In accordance with an embodiment, a similar process can be performed for the other two sources, the difference being that for the OSCS source, the schema may not be known at the outset, so that can instead be obtained by some means (for example, metadata query, sampling, or user defined).

10 **[000189]** In accordance with an embodiment, the sources of data can optionally be profiled to investigate the sources further, which may help in deriving recommendations later in the integration flow.

[000190] In accordance with an embodiment, the next step is to define how the separate sources can be joined together around a central item, which is typically the basis (fact) for the
15 analysis, and which can be accomplished by defining a dataflow pipeline. This can be done directly by creating a pipeline domain-specific language (DSL) script, or by using the guided editor where the user can see the effect on the data at each step and can take advantage of the recommendation service that suggests how the data could be, e.g., corrected, enriched, joined.

20 **[000191]** At this point, the user can request that the system suggest a suitable structure to analyze the resulting content. For example, in accordance with an embodiment, the system can use the knowledge service (functional type classification) to suggest measures and related dimension hierarchies. Once that is completed, the system can recommend the data flows needed to take the blended data from the earlier pipeline and populate the dimensional target
25 structure. Based on the dependency analysis it will also derive and generate the orchestration flow to load / refresh the target schema.

[000192] In accordance with an embodiment, the system can now generate a HUB to host the target structure and associate it via an adapter to a DBCS generating the data definition language (DDL) needed to create the target schema and, for example, deploy the
30 XDML or whatever form BICS can use to generate the models needed to access the newly created schema. This can be populated by executing the orchestration flow and triggering the exhaust service.

[000193] **Figure 5** illustrates an example use of a data flow with a pipeline, in accordance with an embodiment.

35 **[000194]** As illustrated in **Figure 5**, in accordance with an embodiment, the system allows users to define pipelines 302 representative of a data flow, in this example including pipeline

steps S1 through S5, to describe the processing of data when built/executed 304 as applications.

[000195] For example, in accordance with an embodiment, users can invoke ingest, transform, model and publish services, or other services such as, for example, policy 306, execution 310, or persistence services 312, to process data in the pipeline. Users can also define solutions (i.e., a control flow) to specify a unified flow which can integrate related pipelines together. Typically, solutions model a complete use case, for example the loading of a sales cube and associated dimensions.

10 Data AI System Components

[000196] In accordance with an embodiment, adapters enable connection to, and ingest data from, various end points and are specific to application or source type.

[000197] In accordance with an embodiment, the system can include a predefined set of adapters, some of which can leverage other SOA adapters, and allow additional adapters to be registered to the framework. There can be more than one adapter for a given connection type; in which cases, the ingest engine will choose the adapter most suited based on the connection type configuration of the HUB.

[000198] Figure 6 illustrates an example use of an ingest/publish engine and ingest/publish service with a pipeline, in accordance with an embodiment.

20 [000199] As illustrated in Figure 6, in accordance with an embodiment, an ingest / publish engine 330 can be accessed via an ingest / publish service 332 by a pipeline 334, which in this example is designed to ingest data 336 (e.g., sales data) from an input HUB (e.g., SFDC HUB1), transform the ingested data 338, and publish the data to an output HUB 340 (e.g., Oracle HUB).

25 [000200] In accordance with an embodiment, the ingest / publish engine supports a plurality of connection types 331, each of which connection type 342 is associated with one or more adapter(s) 344 that provides access to a HUB.

[000201] For example, as illustrated in the example of Figure 6, in accordance with an embodiment, an SFDC connection type 352 can be associated with an SFDC-Adp1 adapter 354, and an SFDC-Adp2 adapter 356, that provide access to SFDC HUBs 358, 359; while an ExDaaS connection type 362 can be associated with an ExDaaS-Adp adapter 364 that provides access to an ExDaaS HUB 366; and an Oracle Connection type 372 can be associated with an Oracle Adp adapter 374 that provides access to an Oracle HUB 376.

35 Recommendation Engine

[000202] In accordance with an embodiment, the system can include a recommendation

engine or knowledge service that operates as an expert filtering system that predicts / suggests the most relevant from among the several possible actions that can be performed on the data.

[000203] In accordance with an embodiment, recommendations can be chained to facilitate the user to step through them, to achieve the prescribed end goal. For example, the recommendation engine can guide the user through a set of steps in transforming a data set into a data cube to be published to the target BI system.

[000204] In accordance with an embodiment, the recommendation engine utilizes three aspects: (A) business type classification, (B) functional type classification, and (C) a knowledge base. Ontology management and Query/Search functionality on datasets or entities can be provided by, for example, a YAGO3-seeded federated ontology with Query API, MRS and Audit Repository. Business entity classification can be provided by, for example, a ML pipeline based classification to identify the business type. Functional type classification can be provided by, for example, deductive and rule-based functional type classification. Action recommendation can be provided by, for example, inductive and rule-based data prep, transform, model, dependencies and related recommendations.

Classification Service

[000205] In accordance with an embodiment, the system provides a classification service, which can be categorized into business type classification and functional type classification, each of which are further described below.

Business Type Classification

[000206] In accordance with an embodiment, the business type of an entity is its phenotype. The observable characteristics of the individual attributes in the entity are as important as the definition, in identifying the business type of an entity. While the classification algorithm uses the schematic definition of a dataset or entity, it can also utilize the models built using data, to classify the datasets or entities business type.

[000207] For example, in accordance with an embodiment, a dataset ingested from a HUB can be classified as one of existing business types (seeded from the main HUB) that are known to the system, or can be added as a new type if it cannot be classified into an existing one.

[000208] In accordance with an embodiment, business type classification is utilized in making recommendations, either based on inductive reasoning (from transforms defined on similar business type in pipelines), or simple propositions derived from the classification root entity.

[000209] Generally described, in accordance with an embodiment, the following set of

steps describes the classification process: ingest and seeding from a main (training) hub; building models and compute col stats and register them for use in classification; classifying datasets or entities, from newly added hubs, including creating a profile / compute col stats; and classifying datasets or entities, to provide a short list of entity models to use based on structure and col stats; and classifying datasets or entities, including multiclass classification to compute / predict using the models.

5 **[000210]** **Figure 7** illustrates the process of ingest and training from a HUB, in accordance with an embodiment.

10 **[000211]** As illustrated in Figure 7, in accordance with an embodiment, the data from a HUB 382 (e.g., in this example, a RelatedIQ source) can be read by the recommendation engine 380, as datasets 390 (e.g., as Resilient Distributed Datasets, RDDs), including in this example an accounts dataset 391, events dataset 392, contacts dataset 393, lists dataset 394, users dataset 395.

15 **[000212]** In accordance with an embodiment, a plurality of type classification tools 400 can be used with an ML pipeline 402, for example GraphX 404, Wolfram/Yago 406, and/or MLib Statistics 408, are used in seeding the knowledge graph 440 with entity metadata (training or seeding data) when a HUB is first registered.

20 **[000213]** In accordance with an embodiment, the dataset or entity metadata and data are ingested from the source HUB and stored in the data lake. During model generation 410, the entity metadata (attributes and relationship with other entities) is used, for example through FP-growth logistics regression 412, in generating the models 420 and knowledge graph representing all the datasets or entities, in this example representing events 422, accounts 424, contacts 426, and users 428. As part of the seeding, regression models are built using dataset or entity data and attribute statistics (min value, max value, mean, or probability density) are computed.

25 **[000214]** **Figure 8** illustrates the process of building models, in accordance with an embodiment.

30 **[000215]** As illustrated in Figure 8, in accordance with an embodiment, when run in, for example, a Spark environment 430, Spark MLib statistics can be used to compute column stats that are added as attribute properties in the knowledge graph. The computed column stats, along with other dataset or entity metadata, can be used to shortlist the entities whose regression models will be used in testing the new entities for classification.

[000216] **Figure 9** illustrates the process of classifying datasets or entities, from newly added HUBs, in accordance with an embodiment.

35 **[000217]** As illustrated in Figure 9, in accordance with an embodiment, when a new HUB is added, in this example, an Oracle HUB 442, the datasets or entities provided by that HUB,

e.g., party information 444, and customer information 446, are classified by the model as a party 448 based on the training or seeded data created previously.

[000218] For example, in accordance with an embodiment, the column statistics are computed from the data of the new datasets or entities, and a set of predicates representing the entity's sub graph is created using this information along with other metadata available as part of the ingest.

[000219] In accordance with an embodiment, the column statistics computation is useful in maximum likelihood estimation (MLE) methods, while sub graphs as well as in regression models for the dataset. The set of graph predicates generated for the new entity will be used to short list candidate entity models for testing and classifying the new entity.

[000220] **Figure 10** further illustrates the process of classifying datasets or entities, from newly added HUBs, in accordance with an embodiment.

[000221] As illustrated in Figure 10, in accordance with an embodiment, predicates representing a sub graph of new datasets or entities, to be classified, are compared with similar sub graphs representing datasets or entities that are already part of the knowledge graph. A ranking of the matching entities, based on the probability of a match, is used in short-listing the entity models used in testing for classifying new entities.

[000222] **Figure 11** further illustrates the process of classifying datasets or entities, from newly added HUBs, in accordance with an embodiment.

[000223] As illustrated in Figure 11, in accordance with an embodiment, the regression models of the short listed matching datasets or entities are used in testing the data from the new datasets or entities. The ML pipeline can be extended to include additional classification methods/models to improve the accuracy of the process. The classification service will classify the new entity if there exists a match within the acceptable threshold, for example, in this illustration, a probability greater than 0.8. If not, the dataset or entity can be added to the knowledge graph as a new business type. A user can also validate the classification, by accepting or rejecting the result.

Functional Type Classification

[000224] In accordance with an embodiment, the functional type of an entity is its genotype. The functional type can also be described as the interface through which transform actions are defined. For example, a join transform or filter is defined on a functional type, such as a relational entity in this case. In summary, all transforms are defined in terms of functional types as parameters.

[000225] **Figure 12** illustrates an object diagram for use in functional type classification, in accordance with an embodiment.

[000226] As illustrated in Figure 12 by way of object diagram 460, in accordance with an embodiment, the system can describe the general case (in this example, dimension, level, or cube) through a set of rules against which a dataset or entity is evaluated to identify its functional type.

5 **[000227]** For example, in accordance with an embodiment, a multidimensional cube can be described in terms of its measure attributes and dimensionality, each of which can themselves be defined in terms of their types and other characteristics. The rules engine will evaluate a business type entity and annotate its functional type based on the evaluation.

[000228] **Figure 13** illustrates an example of a dimension functional type classification, in accordance with an embodiment.

[000229] As illustrated in the example functional type classification 470 hierarchy shown in Figure 13, in accordance with an embodiment, a level can be defined, for example in terms of its dimensions, and level attributes.

15 **[000230]** **Figure 14** illustrates an example of a cube functional type classification, in accordance with an embodiment.

[000231] As illustrated in the example functional type classification 480 hierarchy shown in Figure 14, in accordance with an embodiment, a cube can be defined, for example, in terms of its measure attributes and dimensionality.

[000232] **Figure 15** illustrates an example usage of functional type classification to evaluate a business entity's functional type, in accordance with an embodiment.

[000233] As illustrated in Figure 15, in this example 490, in accordance with an embodiment, a sales dataset should be evaluated as a cube functional type by the rules engine. Similarly the product, customer, and time should be evaluated as dimensions and levels (for example by age group, gender).

25 **[000234]** In accordance with an embodiment, the rules that identify the entity's functional types and the dataset or entity elements for this particular example are provided below, including several rules that can be specified to evaluate the same functional type. For example, a column of type "Date" can be considered a dimension regardless of whether there is a reference to a parent level entity. Similarly zip code, gender and age may only require
30 data rules to identify them as dimensions:

Customer

Id, Name → (Dimension isComposedOf DimensionAttrib)

AgeGroup, Gender → (Dimension isComposedOf IdAttrib, IdAttrib references

35 Level)

Time

Day → (Dimension /Level isComposedOf DimensionAttrib /LevelAttrib)

Month → (Dimension isComposedOf IdAttrib, IdAttrib references Level)

5 Sales

Qty, Price, Amount → (Cube isComposedOf CubeAttr and Data rule on this columns, for example numeric min/max values, probability density)

Custid → (DimAttr references Dimension, Cube isComposedOf CubeAttr)

Date → (references Dimension, Cube isComposedOf CubeAttr)

10

[000235] Figure 16 illustrates an object diagram for use in functional transformation, in accordance with an embodiment.

[000236] As illustrated in Figure 15, in accordance with this example 500, in accordance with an embodiment, transform functions can be defined on functional types. Business entities (business types) are annotated as a functional type; including by default that complex business types are of functional type "entity".

[000237] Figure 17 illustrates the operation of a recommendation engine, in accordance with an embodiment.

[000238] As illustrated in Figure 17, in accordance with an embodiment, the recommendation engine generates recommendations which are a set of actions defined on a business type. Each action is a directive to apply a transform on a dataset(s).

[000239] In accordance with an embodiment, a recommendation context 530 abstracts the sources of the recommendation, and contains metadata to identify the set of propositions that generated the recommendation. The context allows the recommendation engine to learn and prioritize recommendations based on user's response.

[000240] In accordance with an embodiment, a target entity deducer / mapper 512 uses the definition of the target (and classification service that annotates dataset or entity and attribute business types) to make transform recommendations that facilitate the current dataset to map into the target. This is common when the user starts with a known target object (say sales cube) and builds a pipeline to instantiate the cube.

[000241] In accordance with an embodiment, templates (pipeline/solution) 514 define a reusable set of pipeline steps and transforms to achieve a desired end result. For example, a template might contain steps to enrich, transform, and publish to a data mart. The set of recommendations in this case will reflect the template design.

[000242] In accordance with an embodiment, a classification service 516 identifies the business type of a dataset or entity ingested from a HUB into the data lake. Recommendations

on the entity can be made based on transforms applied on similar entities (business type) or in conjunction with the target entity deducer/mapper.

[000243] In accordance with an embodiment, a functional type service 518 annotates the functional type(s) that a dataset or entity can assume based on rules defined. For example, to generate a cube from a given dataset or join it with dimensional tables, it is important to evaluate if the dataset meets the rules that define the cube functional type.

[000244] In accordance with an embodiment, a pattern inference from pipelines component 520 allows the recommendation engine to summarize transforms performed on a given business type in existing pipeline definitions in similar contexts and suggest similar transform as recommendation in the current context.

[000245] In accordance with an embodiment, the recommendation context can be used to process a recommendation 532, including actions 534, transform function 535, action parameters 536, function parameters 537, and business type 538.

Data Lake / Data Management Strategies

[000246] As described previously, in accordance with an embodiment, a data lake provides a repository for the persistence of information from the system HUB, or other components.

[000247] **Figure 18** illustrates the use of a data lake, in accordance with an embodiment.

[000248] As illustrated in **Figure 18**, in accordance with an embodiment, a data lake can be associated with one or more data access APIs 540, a cache 542, and persistence store 544, that operate together to receive ingested data that has been normalized, for use with a plurality of pipelines 552, 554, 556.

[000249] In accordance with an embodiment, a variety of different data management strategies can be used to manage data (performance, scalability) and its life cycle in the data lake, which can be broadly classified as data-driven or process-driven.

[000250] **Figure 19** illustrates the use of a data-driven strategy to manage a data lake, in accordance with an embodiment.

[000251] As illustrated in **Figure 19**, in accordance with an embodiment, in a data-driven approach, the unit of management is derived based on the HUB or data server definitions. For example, in this approach the data from an Oracle 1 HUB can be stored in a first data center 560 associated with that HUB, while the data from SFHUB1 can be stored in a second data center 562 associated with that HUB.

[000252] **Figure 20** illustrates the use of a process-driven strategy to manage a data lake, in accordance with an embodiment.

[000253] As illustrated in **Figure 20**, in accordance with an embodiment, in a process-

driven approach, the unit of management is derived based on the related pipelines accessing the data. For example, in this approach the data associated with a Sales pipeline can be stored in a first data center 564 associated with that pipeline, while the data from other pipelines (e.g., pipelines 1, 2, 3) can be stored in a second data center 566 associated with those other pipelines.

Pipelines

[000254] In accordance with an embodiment, a pipeline defines the transformation or processing that is to be performed on the ingested data. The processed data can be stored in the data lake or can be published to another end point, such as, for example, DBCS.

[000255] **Figure 21** illustrates the use of a pipeline compiler, in accordance with an embodiment.

[000256] As illustrated in **Figure 21**, in accordance with an embodiment, a pipeline compiler 582 operates between design 570 and execution 580 environments, including accepting one or more pipeline metadata 572, and a DSL, e.g., Java DSL 574, JSON DSL 576, Scala DSL 578, and providing an output for use with the execution environment e.g., as a Spark application 584 and/or SQL statements 586.

[000257] **Figure 22** illustrates an example pipeline graph, in accordance with an embodiment.

[000258] As illustrated in **Figure 22**, in accordance with an embodiment, a pipeline 588 comprises a list of pipeline steps. Different types of pipeline steps represent different kinds of operations that can be performed in the pipeline. Each pipeline step can have a number of input data sets and a number of output data sets, generally described by pipeline step parameters. The processing order of operations in the pipeline is defined by binding the output pipeline step parameters from a preceding pipeline step to a subsequent pipeline step. In this manner, pipeline steps and the relationships between pipeline step parameters form a directed acyclic graph (DAG).

[000259] In accordance with an embodiment, a pipeline can be reused in another pipeline, if the pipeline contains one or more special pipeline steps (signature pipelines) that represent the input and output pipeline step parameters of the pipeline. The enclosing pipeline refers to the pipeline to be reused through a (pipeline usage) pipeline step.

[000260] **Figure 23** illustrates an example of a data pipeline, in accordance with an embodiment.

[000261] As illustrated in the example data pipeline 600 shown in **Figure 23**, in accordance with an embodiment, data pipelines perform data transformations. The data flow in the pipeline is represented as a binding of pipeline step parameters. Various types of

pipeline steps are supported for different transformation operations, including, for example: Entity (retrieve data from data lake or publish processed data to data lake/other HUBs); and Join (fusion of multiple sources).

[000262] Figure 24 illustrates another example of a data pipeline, in accordance with an embodiment.

[000263] As illustrated in the example data pipeline 610 shown in Figure 24, in accordance with an embodiment, a data pipeline P1 can be reused in another data pipeline P2.

[000264] Figure 25 illustrates an example of an orchestration pipeline, in accordance with an embodiment.

[000265] As illustrated in the example orchestration pipeline 620 shown in Figure 25, in accordance with an embodiment, using an orchestration pipeline, pipeline steps can be used to represent the task or job that needs to be executed in an overall orchestration flow. All of the pipeline steps in an orchestration pipeline are assumed to have one input pipeline step parameter and one output pipeline step parameter. The execution dependency between tasks can be expressed as the binding between pipeline step parameters.

[000266] In accordance with an embodiment, parallel execution of tasks can be scheduled if the pipeline steps are dependent on the same preceding pipeline step without conditions (i.e. a fork). If a pipeline step is dependent on multiple preceding paths, then the pipeline step will wait for all the multiple paths to be completed before its own execution (i.e., a join). However, this does not always imply the tasks are executed in parallel. The orchestration engine can decide whether to execute the tasks in serial or in parallel depending on available resources

[000267] In the example illustrated in Figure 25, in accordance with an embodiment, a pipeline step 1 is first executed. If a pipeline step 2 and a pipeline step 3 are executed in parallel, then a pipeline step 4 will be executed only when pipeline step 2 and pipeline step 3 are both finished. The orchestration engine can also execute this orchestration pipeline serially as (pipeline step 1, pipeline step 2, pipeline step 3, pipeline step 4), or (pipeline step 1, pipeline step 3, pipeline step 2, pipeline step 4), as long as it satisfies the dependency between pipeline steps.

[000268] Figure 26 further illustrates an example of an orchestration pipeline, in accordance with an embodiment.

[000269] As illustrated in the example pipeline 625 shown in Figure 26, in accordance with an embodiment, each pipeline step can return a status 630, such as, for example, a success or error status depending on its own semantics. The dependency between two pipeline steps can be conditional, based on the return status of the pipeline step. In the

example shown, pipeline step 1 is first executed; if it completes successfully, then pipeline step 2 will be executed, otherwise pipeline step 3 will be executed. After either pipeline step 2 or pipeline step 3 is executed, then pipeline step 4 will be executed.

[000270] In accordance with an embodiment, orchestration pipelines can be nested, so that one orchestration pipeline can refer to another orchestration pipeline through a pipeline usage. An orchestration pipeline can also refer to a data pipeline as a pipeline usage. The difference between an orchestration pipeline and a data pipeline is that orchestration pipeline refers to a data pipeline that does not contain signature pipeline steps, while a data pipeline can reuse another data pipeline that contains signature pipeline steps.

[000271] In accordance with an embodiment, depending on the types of pipeline steps and code optimization, a data pipeline can be generated as a single Spark application for execution in a Spark cluster, as multiple SQL statements for execution in DBCS, or as a mixture of SQL and Spark code. For an orchestration pipeline, it can be generated for execution in the underlying execution engine, or in a workflow schedule component such as, e.g., Oozie.

Coordination Fabric

[000272] In accordance with an embodiment, a coordination fabric or fabric controller provides the necessary tools to deploy and manage framework components (service provider) and applications (user designed), manages application execution and resource requirements/allocation, and provides a integration framework (messaging bus) to facilitate communication between the various components.

[000273] **Figure 27** illustrates the use of a coordination fabric including a messaging system, in accordance with an embodiment.

[000274] As illustrated in **Figure 27**, in accordance with an embodiment, a messaging system (e.g., Kafka) 650 coordinated interactions between a resource manager 660 (e.g., Yarn/Mesos), scheduler 662 (e.g., Chronos), application schedulers 664 (e.g., Spark), and a plurality of nodes, here indicated as nodes 652, 654, 656, 658.

[000275] In accordance with an embodiment, the resource manager is used to manage the life cycle of data computation tasks/applications including scheduling, monitoring, application execution, resource arbitration and allocation, load balancing; including managing the deployment and configuration of components (that are producers and consumers of messages) in a message-driven component integration framework; upgrade of components (services) without downtime; and upgrade of infrastructure with minimal or no disruption of service.

[000276] **Figure 28** further illustrates the use of a coordination fabric including a

messaging system, in accordance with an embodiment.

[000277] As illustrated in Figure 28, in accordance with an embodiment, the dependencies across components in the coordination fabric are illustrated through a simple data-driven pipeline execution use case, where (c) indicates a consumer, and (p) indicates a
5 producer.

[000278] In accordance with the embodiment illustrated in Figure 28, a scheduler (p) starts the process (1) by initiating ingest of data to a HUB. An ingest engine(c) processes the request (2) ingesting data from the HUB into the data lake. After the ingest process is completed, the ingest engine (p) communicates (3) the completion status to initiate pipeline
10 processing. If the scheduler supports data-driven execution, it can automatically initiate (3a) the pipeline process to execute. The pipeline engine (c) computes (4) the pipelines that are waiting on the data for execution. The pipeline engine (p) communicates (5) the list of pipeline applications to schedule for execution. The scheduler gets the execution schedule request (6) for the pipelines; and initiates execution (6a) of the pipelines. The application scheduler (e.g.,
15 Spark) arbitrates (7) with the resource manager for resource allocation to execute the pipelines. The application scheduler sends the pipeline (8), for execution to the executor in the allocated nodes.

On-Premise Agent

[000279] In accordance with an embodiment, an on-premise agent facilitates access to
20 local data and in a limited fashion, distributed pipeline execution. The on-premise agent is provisioned and configured to communicate with, e.g., a Cloud DI service, to process data access and remote pipeline execution requests.

[000280] **Figure 29** illustrates an on-premise agent for use with a system, in accordance
25 with an embodiment.

[000281] As illustrated in Figure 29, in accordance with an embodiment, a cloud agent adapter 682 provisions (1) the on-premise agent 680 and configures the agent adapter end point for communication.

[000282] An ingest service initiates (2) the local data access request for HUB1 through
30 the messaging system. a cloud agent adapter operates as an intermediary (3) between on-premise agent and the messaging system by providing access to requests initiated through the ingest service as well as writing data from on-premise agent into the data lake and notifying the completion of the task through messaging system.

[000283] A premise agent polls (4) the cloud agent adapter for data access requests to
35 process or upload data into the cloud. The cloud agent adapter writes data (5) into the data lake and notifies pipelines through the messaging system.

DFML Flow Process

[000284] Figure 30 illustrates a data flow process, in accordance with an embodiment.

[000285] As illustrate in Figure 30, in accordance with an embodiment, during an ingest step 692, data is ingested from various sources, for example, SFDC, S3, or DBaaS.

5 [000286] During a data prep step 693, the ingested data is prepared for example by being de-duplicated, standardized, or enriched.

[000287] During a transform step 694, the system performs merges, filters, or lookups at datasets, to transform the data.

[000288] During a model step 695, one or more models are generated, together with
10 mappings to the models.

[000289] During a publish step 696, the system can publish models, specify policies and schedules, and populate target data structures.

Metadata and Data-driven Auto-Mapping

15 [000290] In accordance with an embodiment, the system can provide support for auto-mapping of complex data structures, datasets or entities, between one or more sources or targets of data (referred to herein in some embodiments as HUBs). The auto-mapping can be driven by a metadata, schema, and statistical profiling of a dataset; and used to map a source dataset or entity associated with an input HUB, to a target dataset or entity, or vice versa, to
20 produce an output data prepared in a format or organization (projection) for use with one or more output HUBs.

[000291] For example, in accordance with an embodiment, for a user implementing (e.g., building) a dataflow, pipeline, or Lambda application, the user may desire to choose data to be mapped from a source or input dataset or entity, within an input HUB, to a target or output
25 dataset or entity, within an output HUB.

[000292] In accordance with an embodiment, since generating a map of data from an input HUB to an output HUB for very large set of HUBs and dataset or entities by hand can be an extremely time consuming and inefficient task, auto-mapping can enable a user to focus on simplification of a dataflow application, e.g., pipeline, Lambda application, by providing a user
30 with recommendations for mapping data.

[000293] In accordance with an embodiment, the data AI subsystem can receive an auto-map request for an auto-map service via a graphical user interface (e.g., a Lambda Studio Integrated Development Environment (IDE)).

[000294] In accordance with an embodiment, the request can include a file specified for
35 an application for which the auto-map service is to be performed, together with information identifying an input HUB, a dataset or entity, and one or more attributes. The application file

can include information about data for the application. The data AI subsystem can process the application file to extract entity name and other shape characteristics of the entity including attribute names and data type, which the auto-map service can use in searching to find a potential candidate set for mapping.

5 **[000295]** In accordance with an embodiment, the system can access data for transformation into a HUB, such as, for example, a data warehouse. The accessed data can include various types of data including semi-structured and structured data. The data AI subsystem can perform a metadata analysis on the accessed data, including determining one or more shapes, features, or structure of the data. For example, the metadata analysis can
10 determine the types of the data (e.g., business types and functional types), and a columnar shape of the data.

[000296] In accordance with an embodiment, based on the metadata analysis of the data, one or more samples of data can be identified, and a machine learning process applied to the sampled data, to determine a category of data in the accessed data, and update a model.
15 The category of data, for example, may indicate the relevant portions of the data, such as a fact table in the data.

[000297] In accordance with an embodiment, the machine learning can be implemented using, for example, a logistic regression model, or other type of machine learning model that can be implemented for machine learning. In accordance with an embodiment, the data AI
20 subsystem can analyze a relationship of one or more data items in the data, based on the category of data, the relationship indicating one or more fields in the data for the category of data.

[000298] In accordance with an embodiment, the data AI subsystem can perform a process for feature extraction, including determining one or more metadata, data type, and
25 statistical profiles of randomly sampled data for attributes of the accessed data.

[000299] For example, in accordance with an embodiment, the data AI subsystem can generate a profile of the accessed data based on its category of data. The profile can be generated for transformation of the data into an output HUB, and, for example, displayed in a graphical user interface.

30 **[000300]** In accordance with an embodiment, as a result of creating such profiles, the model can support a recommendation, with a degree of confidence, for how a candidate dataset or entity is similar to an input dataset or entity. The recommendations can be filtered and sorted and then provided to the user via the graphical user interface.

[000301] In accordance with an embodiment, an auto-map service can dynamically
35 suggest recommendations based on a stage in which a user is building a dataflow application, e.g., pipeline, Lambda application.

[000302] An example of a recommendation at an entity level may include, in accordance with an embodiment, a recommendation for an attribute, for example a column of the entity to be automatically mapped to another attribute, or another entity. The service can continuously provide recommendations and guide the user based on user's previous activities.

5 **[000303]** In accordance with an embodiment, the recommendations can be mapped from, e.g., a source dataset or entity associated with an input HUB, to a target dataset or entity associated with an output HUB, using an application programming interface (API) (e.g., a REST API) provided by the auto-map service. The recommendation can indicate a projection of data, such as, for example, an attribute, a data type, and an expression, wherein the
10 expression can be a mapping of the attribute for the data type.

[000304] In accordance with an embodiment, the system can provide a graphical user interface to select an output HUB for transformation of the accessed data based on the recommendation. For example, the graphical user interface can enable a user to select a recommendation for transformation of data to an output HUB.

15

Auto-Mapping

[000305] In accordance with an embodiment, the auto-mapping functionality can be defined mathematically, in which an entity set E is defined as:

$$20 \quad E = \{e_1, e_2, \dots, e_n \mid \forall e_i \in S\}$$

$$\text{Shape: } S = \{\text{MetaData} \times \text{DataType} \times \text{StatisticalProfile}\}$$

[000306] Where the shape set S includes metadata, data type and statistical profiling
25 dimensions. The goal is to find j such that the probability of similarity between e_i and e_j is the highest.

$$e_j^* = \underset{j}{\text{argmax}} p\{\text{Sim}(e_i, e_j) \mid \text{entity} = e_i\}$$

30 **[000307]** At dataset or entity level, the problem is a binary problem, i.e., whether the dataset or entity is similar or dissimilar. Let $f_s, f_t, h(f_s, f_t)$ denote sets of features for source, target, and interactive features between source and target. Therefore the goal is to estimate the probability of similarity:

$$35 \quad p = g(f_s, f_t, h(f_s, f_t); \beta)$$

$$g(\cdot): [0,1]^q \rightarrow [0,1]$$

[000308] The log-likelihood function is defined as:

$$\ell(\boldsymbol{\beta}) = \sum_{q=1}^Q c_q \log p_q(\boldsymbol{\beta}) + (1 - c_q) \log(1 - p_q(\boldsymbol{\beta}))$$

5 [000309] Therefore, in the logistic regression model the unknown coefficients can be estimated as follows:

$$g(x; \boldsymbol{\beta}) = \frac{1}{1 + e^{-\boldsymbol{\beta}^T x}}$$

10 $\boldsymbol{\beta}^* = \arg_{\boldsymbol{\beta}} \max \ell(\boldsymbol{\beta})$

[000310] In accordance with an embodiment, the auto-map service can be triggered by, for example, receiving a HTTP POST request from the system facade service. The system facade API passes the dataflow application, e.g., pipelines, Lambda application JSON file from
 15 UI to the auto-map REST API, and the parser module processes the application JSON file and extracts entity name and shape of dataset or entity including attribute names and data type.

[000311] In accordance with an embodiment, the auto-map service uses a search to find quickly a potential candidate set for mapping. The candidate set needs to be a highly relevant set, and therefore a special index and query can be used to achieve this. This special index
 20 incorporates a special search field where all attributes of an entity are stored and tokenized with all N-gram combinations. At query time, the search query builder module leverages a fuzzy search feature based, for example, on Levenshtein distance, to construct a special query using both entity name and attribute names of the given entity and leverage the search boost function to sort results by their relevance in sense of string similarity.

25 [000312] In accordance with an embodiment, the recommendation engine shows a number of relevant results to users, for example a selection of top N results in most cases.

[000313] In accordance with an embodiment, in order to achieve a high precision, a machine learning model compares pairs of source and targets and score similarity of entities based on extracted features. The feature extraction includes metadata, data type and
 30 statistical profiles of randomly sampled data for each attributes.

[000314] Although the descriptions provided here, in accordance with an embodiment, generally describes the use of a logistic regression model to learn auto-mapping examples taken from an Oracle Business Intelligence (OBI) lineage mapping data, other supervised machine learning models can be used instead.

35 [000315] In accordance with an embodiment, the output of the logistic regression model

represents an overall confidence of how a candidate dataset or entity is similar to an input dataset or entity in a statistical sense. In order to find an exact mapping one or more other models can be used to calculate similarity of source attribute with target attribute using similar features.

5 **[000316]** Finally, in accordance with an embodiment, the recommendations are filtered and sorted and sent back to the system facade and passed to the user interface. The auto-map service dynamically suggests recommendations based on at what stage the user is during the dataflow application, e.g., pipeline or Lambda application design. The service can continuously provide recommendations and guide the user, based on the user's previous
10 activities. The auto-mapping can be performed in either forward engineering or reverse engineering senses.

[000317] **Figure 31** illustrates an automated mapping of data types, in accordance with an embodiment.

[000318] As illustrated in Figure 31, in accordance with an embodiment, a system façade
15 701 and auto-map API 702 allows for the receipt of a dataflow application, e.g., pipeline or Lambda application from the software development component, e.g., Lambda Studio. A parser 704 processes the application's JSON file, and extracts entity names and shapes, including attribute names and data type.

[000319] In accordance with an embodiment, a search index 708, is used to support a
20 primary search 710, to find a potential candidate set of datasets or entities, for mapping. A search query builder module 706 constructs a query using both entity name and attribute names of the given entity, to determine a selection of datasets or entities 712.

[000320] In accordance with an embodiment, a machine learning (ML) model is used to
25 compare pairs of sources and targets, and to score similarity of datasets or entities, based on extracted features. The feature extraction 714 includes metadata, data type and statistical profiles of randomly sampled data for each attributes.

[000321] In accordance with an embodiment, a logistic regression model 716 provides,
30 as an output, an overall confidence of how a candidate entity is similar to an input entity. In order to find a more exact mapping, a column mapping model 718 is used to further evaluate similarity of source attribute with target attribute.

[000322] In accordance with an embodiment, the recommendations are then sorted as
35 auto-mappings 720, for return to the software development component, e.g, Lambda Studio. The auto-map service dynamically suggests recommendations based on at which stage the user is, during the dataflow application, e.g., pipeline or Lambda application design. The service can continuously provide recommendations and guide the user, based on the user's previous activities.

[000323] **Figure 32** illustrates an auto-map service for generation of mappings, in accordance with an embodiment.

[000324] As illustrated in Figure 32, in accordance with an embodiment, an auto-map service can be provided for generation of mappings, including wherein a UI query 728 is received and passed to a query understanding engine 729, and then to a query decomposition 730 component.

[000325] In accordance with an embodiment, a primary search 710 is performed using a data HUB 722 to determine candidate datasets or entities 731, for use in subsequent metadata and stats profile processing 732.

[000326] In accordance with an embodiment, the result are passed to a get stats profile 734 component and the data AI system 724 provide feature extraction 735. Results are used for synthesis 736, final confidence merging and ranking 739 according to the models 723, and providing of recommendations and associated confidences 740.

Auto-Map Examples

[000327] **Figure 33** illustrates an example of mapping between a source schema and target schema, in accordance with an embodiment.

[000328] As illustrated in Figure 33, the example 741 shows, in accordance with an embodiment, a simple auto-mapping example based on, for example, (a) hypernyms, (b) synonyms, (c) equality, (d) Soundex, and (e) fuzzy matching.

[000329] **Figure 34** illustrates another example of mapping between a source schema and target schema, in accordance with an embodiment.

[000330] As illustrated in Figure 34, an approach solely based on metadata will fail if this information is irrelevant. In accordance with an embodiment, Figure 34 illustrates an example 742 wherein the source and target attribute names are completely uninformative. When there is a lack of metadata features, the system can employ models that include statistical profiling of features to achieve finding similar entities.

Auto-Map Process

[000331] **Figure 35** illustrates a process for providing automated mapping of data types, in accordance with an embodiment.

[000332] As illustrated in Figure 35, at step 744, in accordance with an embodiment, an accessed data is processed, to perform a metadata analysis of the accessed data.

[000333] At step 745, one or more samples of the accessed data are identified.

[000334] At step 746, a machine learning process is applied to determine a category of data within the accessed data.

[000335] At step 748, a profile of the accessed data is generated, based on the determined category of data, for use in auto-mapping the accessed data.

Dynamic Recommendations and Simulation

5 [000336] In accordance with an embodiment, the system can include a software development component (referred to herein in some embodiments as Lambda Studio), and graphical user interface (referred to herein in some embodiments as pipeline editor, or Lambda Studio IDE), that provides a visual environment for use with the system, including providing real-time recommendations for performing semantic actions on data accessed from an input
10 HUB, based on an understanding of the meaning or semantics associated with the data.

[000337] For example, in accordance with an embodiment, the graphical user interface can provide real-time recommendations for performing operations (also referred to as semantic actions) on data accessed from an input HUB, including a partial data, shape or other characteristics of that data. The semantic actions can be performed on the data based on the
15 meaning, or semantics associated with data. The meaning of the data can be used to select semantic actions that can be performed on the data.

[000338] In accordance with an embodiment, a semantic action may represent an operator on one or more datasets, and can reference a base semantic action or function declaratively defined in the system. One or more processed datasets can be generated by
20 performing a semantic action. A semantic action can be defined by parameters that are associated with specific functional or business types. They represent a specific upstream dataset to be processed. A graphical user interface can be metadata-driven, such that the graphical user interface is generated dynamically to provide recommendations based on metadata identified in the data.

25 [000339] **Figure 36** illustrates a system that displays one or more semantic actions enabled for accessed data, in accordance with an embodiment.

[000340] As illustrated in Figure 36, in accordance with an embodiment, using a graphical user interface 750 having a user input area 752 a query for semantic actions enabled for the accessed data is sent to a knowledge source of the system, wherein the query indicates the
30 classification of the accessed data.

[000341] In accordance with an embodiment, a response to the query is received from the knowledge source, wherein the response indicates one or more semantic actions enabled for the accessed data and identified based on the classification of the data.

[000342] In accordance with an embodiment, selected ones of the semantic actions
35 enabled for the accessed data are displayed, for selection and use with the accessed data, including automatically providing or updating a list of the selected ones of the semantic actions

756, or recommendations 758, enabled for the accessed data, during the processing of the accessed data.

[000343] In accordance with an embodiment, the recommendations can be provided dynamically rather than being pre-computed based on static data. For example, the system can provide recommendations in real-time based on data accessed in real-time, taking into account information such as, for example, a user profile or the user's experience level. The recommendations provided by the system for real-time data may be salient, relevant, and precise for producing dataflow applications, e.g., pipelines, Lambda applications. The recommendations can be provided based on user behavior with respect to data associated with particular metadata. The system can recommend semantic actions on information.

[000344] For example, in accordance with an embodiment, the system can ingest data, transform, integrate and publish the data to an arbitrary system. The system can recommend that an entity be used to analyze some of its numerical measures in an interesting analytic manner; pivoting that data on various dimensions, indicating what are the interesting dimensions even, summarizing the data for those dimension hierarchies and enriching the data with more insight.

[000345] In accordance with an embodiment, the recommendations can be provided based on analysis of data using techniques, such as, for example, a metadata analysis of the data.

[000346] In accordance with an embodiment, the metadata analysis can include determining a classification of the data, such as, for example, shapes, features, and structure of the data. The metadata analysis can determine the types of the data (e.g., business types and functional types). The metadata analysis can also indicate a columnar shape of the data. In accordance with an embodiment, the data can be compared to the metadata structure (e.g., shapes and features) to determine a type of data and attributes associated with the data. The metadata structure can be defined in a system HUB (e.g., a knowledge source) of the system.

[000347] In accordance with an embodiment, using the metadata analysis, the system can query a system HUB to identify semantic actions based on the metadata. Recommendations can be semantic actions determined based on analysis of metadata of data accessed from an input HUB. Specifically, semantic actions can be mapped to metadata. For example, semantic actions can be mapped to metadata for which those actions are permitted and/or applicable. The semantic actions can be user defined, and/or can be defined based on a structure of data.

[000348] In accordance with an embodiment, the semantic actions can be defined based on conditions associated with metadata. The system HUB can be modified such that semantic actions are modified, deleted, or augmented.

[000349] Examples of semantic actions may include, in accordance with an embodiment, building a cube, filtering data, grouping data, aggregating data, or other actions that can be performed on data. By defining semantic actions based on metadata, no mapping or scheme may be needed to determine semantic actions permitted for data. Semantic actions can be defined as new and different metadata structures are discovered. As such, the system can dynamically determine recommendations based on identification of semantic actions using metadata analyzed for data received as input.

[000350] In accordance with an embodiment, semantic actions can be defined by a third party, such that the third party can supply data, such as, for example, data defining one or more semantic actions associated with metadata. The system can dynamically query the system HUB to determine the semantic actions available for metadata. As such, the system HUB can be modified, such that the system determines the semantic actions permitted at the time based on such modifications. The system can perform operations (e.g., filter, detect, and register) to process data obtained from a third party, where the data defines semantic actions; and can make semantic actions available based on the semantic actions identified by the processing.

[000351] **Figure 37** and **Figure 38** illustrate a graphical user interface that displays one or more semantic actions enabled for accessed data, in accordance with an embodiment.

[000352] As illustrated in **Figure 37**, in accordance with an embodiment, a software development component (e.g., Lambda Studio), can provide a graphical user interface (e.g., a pipeline editor, or Lambda studio IDE) 750, that can display recommended semantic actions, for use in processing an input data, or simulating the processing of the input data, for projection onto an output HUB.

[000353] For example, in accordance with an embodiment, the interface in **Figure 37** allows a user to display options 752 associated with a dataflow application, e.g., pipeline, Lambda application, including, for example, an input HUB definition 754.

[000354] In accordance with an embodiment, during the creation of a dataflow application, e.g., pipeline, Lambda application, or the simulation of a dataflow application, e.g., pipeline, Lambda application, with an input data, one or more semantic actions 756, or other recommendations 758, can be displayed on the graphical user interface, for review by the user.

[000355] In accordance with an embodiment, in a simulation mode, the software development component (e.g., Lambda Studio) provides a sandbox-environment that allows the user to immediately see the results of performing various semantic actions on the output, including automatically updating a list of semantic actions appropriate to the accessed data, during the processing of the accessed data.

[000356] For example, as illustrated in **Figure 38**, in accordance with an embodiment,

from a starting point of the user searching for some information, the system can recommend operations 760 on information, for example that an entity be used to analyze some of its numerical measures in an interesting analytic manner; pivoting that data on various dimensions, indicating what are the interesting dimensions even, summarizing the data for those dimension hierarchies and also enriching the data with more insight.

5 [000357] In accordance with an embodiment, in the example illustrated, both sources and dimensions have been recommended for an analyzable entity in the system, making the task of building a multi-dimensional cube one of almost point and click.

10 [000358] Typically, such activities require a lot of experience and domain specific knowledge. Using machine learning to analyze both the data characteristics and the user's behavioral patterns for common integration patterns, together with the combination of semantic search and recommendations from machine learning, allows for an approach that offers state of the art tooling for application development for building business focused applications.

15 [000359] **Figure 39** illustrates a process for displaying one or more semantic actions enabled for accessed data, in accordance with an embodiment.

[000360] As illustrated in Figure 39, at step 772, in accordance with an embodiment, an accessed data is processed, to perform a metadata analysis of the accessed data, wherein the metadata analysis includes determining a classification of the accessed data.

20 [000361] At step 774, a query for semantic actions enabled for the accessed data is sent to a knowledge source of the system, wherein the query indicates the classification of the accessed data.

[000362] At step 775, a response to the query is received from the knowledge source, wherein the response indicates one or more semantic actions enabled for the accessed data and identified based on the classification of the data.

25 [000363] At step 776, at a graphical user interface, selected ones of the semantic actions enabled for the accessed data are displayed, for selection and use with the accessed data, including automatically providing or updating a list of the selected ones of the semantic actions enabled for the accessed data, during the processing of the accessed data.

30 **Functional Decomposition of Data Flows**

[000364] In accordance with an embodiment, the system can provide a service to recommend actions and transformations, on an input data, based on patterns identified from the functional decomposition of a data flow for a software application, including determining possible transformations of the data flow in subsequent applications. Data flows can be decomposed into a model describing transformations of data, predicates, and business rules applied to the data, and attributes used in the data flows.

35

[000365] Figure 40 illustrates support for evaluation of a pipeline, Lambda application, into its constituent parts, to facilitate pattern detection and inductive learning, in accordance with an embodiment.

[000366] As illustrated in Figure 40, in accordance with an embodiment, a functional decomposition logic 800 or software component, which can be provided as software or program code that is executable by a computer system or other processing device, and can be used to provide a functional decomposition 802 and recommendations 804, for display 805 (for example, within a pipeline editor, or Lambda Studio IDE). For example, the system can provide a service to recommended actions and transformations on data based on patterns / templates identified from functional decomposition of data flows for dataflow applications, e.g., pipelines, Lambda applications, i.e., observe, through functional decomposition of data flows, patterns for determining possible transformations of data flow in subsequent applications.

[000367] In accordance with an embodiment, the service can be implemented by a framework that can decompose, or break down, data flows into a model describing transformations of data, predicates, and business rules applied to the data, and attributes used in the data flows.

[000368] Traditionally, a data flow for applications may represent a series of transformations on data and the type of transformations applied to data is highly contextual. In most data integration frameworks, process lineage is usually limited or nonexistent in how data flows are persisted, analyzed and generated. In accordance with an embodiment, the system enables deriving contextually relevant patterns from flows or graphs based on semantically rich entity types and to further learn data flow grammar and models and using it to generate complex data flow graphs given similar contexts.

[000369] In accordance with an embodiment, the system can generate one or more data structures defining patterns and templates based on design specification of data flows. Data flows can be decomposed into data structures defining functional expressions to determine patterns and templates. The data flows can be used to predict and generate functional expressions for determining patterns for recommendations of data transformations, wherein the recommendations are based on the models derived from inductive learning of decomposed data flows and inherent patterns, and can be finely grained (for example, recommending a scalar transformation on a particular attribute or the use of one or more attributes in a predicate for filter or join).

[000370] In accordance with an embodiment, dataflow applications, e.g., pipelines, Lambda applications can enable users to generate complex data transformations based on semantic actions on data. The system can store data transformations as one or more data structures defining the flow of data for a pipeline, Lambda application.

[000371] In accordance with an embodiment, decomposing data flow for a dataflow application, e.g., pipeline, Lambda application, can be used to determine pattern analysis of the data and to generate functional expressions. The decomposition can be performed for semantic actions as well as transformations and predicates, or business rules. Each of previous applications semantic actions can be identified through decomposition. Using a process of induction, business logic can be extracted from a data flow including its context elements (business types and functional types).

[000372] In accordance with an embodiment, a model can be generated for the process; and, based on induction, prescriptive data flow design recommendations can be generated which are contextually rich. The recommendations can be based on patterns inference from the model, wherein each of the recommendations may correspond to a semantic action that can be performed on data for an application.

[000373] In accordance with an embodiment, the system can perform a process for inferring patterns for data transformation based on functional decomposition. The system can access a data flow for one or more dataflow applications, e.g., pipelines, Lambda applications. The data flow can be processed to determine one or more functional expressions. The functional expressions can be generated based on actions, predicates, or business rules identified in the data flow. The actions, predicates, or business rules can be used to identify (e.g., infer) the pattern of transformation on the data flow. Inferring a pattern of transformation may be a passive process.

[000374] In accordance with an embodiment, the pattern of transformation can be determined in a crowd sourcing manner based on passive analysis of data flows for different applications. The pattern can be determined using machine learning (e.g., deep reinforcement learning).

[000375] In accordance with an embodiment, a pattern of transformation can be identified for the functions expressions generated for a dataflow application, e.g., pipeline, Lambda application. One or more data flows can be decomposed to infer patterns for data transformation.

[000376] In accordance with an embodiment, using the pattern, the system can recommend one or more data transformations for a data flow of the new dataflow application, e.g., pipeline, Lambda application. In an example of a data flow of processing on data for a monetary exchange, the system can identify a pattern of transformations on the data. The system can also recommend one or more transformations for a new data flow of an application, wherein the data flow involves data for a similar monetary exchange. The transformation(s) can be performed in a similar manner according to the pattern such that the new data flow is modified according to the transformation(s) to produce a similar monetary exchange.

[000377] **Figure 41** illustrates a means of identifying a pattern of transformation in a data flow, for one or more functional expressions generated for each of one or more applications, in accordance with an embodiment.

[000378] As described previously, in accordance with an embodiment, pipelines, e.g., Lambda applications, allow users to define complex data transformations based on semantic actions that correspond to operators in relational calculus. The data transformations are usually persisted as directed acyclic graphs or queries, or, in the case of DFML, as nested functions. Decomposing and serializing a dataflow application, e.g., pipeline, Lambda application, as a nested function enables pattern analysis of the data flow and induce a data flow model that can then be used to generate functional expressions abstracting complex transformations on datasets in similar contexts.

[000379] In accordance with an embodiment, the nested function decomposition is performed not only at the level of semantic actions (row or dataset operators), but also at the scalar transformations and predicate structures, which allows deep lineage capability of a complex data flow. The recommendations based on the induced models can be finely grained (for example, recommending a scalar transformation on a particular attribute or the use of one or more attributes in a predicate for filter or join).

[000380] In accordance with an embodiment, the elements of the functional decomposition generally include that:

[000381] An application represents a top level data flow transformation.

[000382] An action represents an operator on one or more datasets (dataframes to be specific).

[000383] Actions reference the base semantic action or function declaratively defined in the system. Actions can have one or more action parameters each of which can have specific role (in, out, in/out) and type, return one or more processed datasets, and can be embedded or nested several levels deep.

[000384] Action parameters are owned by actions and have specific functional or business types, and represent a specific upstream dataset to be processed. Binding parameters represent datasets or entities in a HUB that are used in the transformation. Value parameters represent intermediate or transient data structures that are processed in the context of the current transformation

[000385] Scope resolvers allow the derivation of process lineage for a dataset or an element in the dataset used in the entire data flow.

[000386] **Figure 42** illustrates an object diagram for use in identifying a pattern of transformation in a data flow, for one or more functional expressions generated for each of one or more applications, in accordance with an embodiment.

[000387] As illustrated in Figure 42, in accordance with an embodiment, a functional decomposition logic can be used to decompose, or break down, the data flow for a dataflow application, e.g., pipeline, Lambda application, into a model describing transformations of data, predicates, and business rules applied to the data, and attributes used in the data flows into a selection of registries, including for example a pattern, or a template 812 (if a template is associated with the pipeline, Lambda application), a service 814, a function 816, function parameters 818, and function type 820.

[000388] In accordance with an embodiment, each of these functional component can be further decomposed into, for example, tasks, 822, or actions 824, reflecting the dataflow application, e.g., pipeline, Lambda application.

[000389] In accordance with an embodiment, a scope resolver 826 can be used to resolve the reference to a particular attribute or embedded object, through its scope. For example, as illustrated in Figure 42, the scope resolver resolves the reference to an attribute or embedded object through its immediate scope. For example, a join function using the output of a filter and another table would have references to both as its scope resolver, and can be used in combination with an InScopeOf operation to resolve leaf node to its root node path.

[000390] Figure 43 illustrates a process of identifying a pattern of transformation in a data flow, for one or more functional expressions generated for each of one or more applications, in accordance with an embodiment.

[000391] As illustrated in Figure 43, in accordance with an embodiment, at step 842, a data flow is accessed for each of one or more software applications.

[000392] At step 844, the data flow for the one or more software applications is processed to generate one or more functional expressions representing the data flow, wherein the one or more functional expressions are generated based on semantic actions and business rules identified in the data flow.

[000393] At step 845, a pattern of transformation in the data flow is identified, for the one or more functional expressions generated for each of the one or more software applications, wherein the semantic actions and business rules are used to identify the pattern of transformation in the data flow.

[000394] At step 847, using the pattern of transformation identified in the data flow, a recommendation of one or more data transformations is provided for a data flow of another software application.

Ontology Learning

[000395] In accordance with an embodiment, the system can perform an ontology analysis of a schema definition, to determine the types of data and datasets or entities,

associated with that schema; and generate, or update, a model from a reference schema that includes an ontology defined based on relationships between entities and their attributes. A reference HUB including one or more schemas can be used to analyze data flows, and further classify or make recommendations for, e.g., transformations enrichments, filtering, or cross-entity data fusion of an input data.

[000396] In accordance with an embodiment, the system can perform an ontology analysis of a schema definition to determine an ontology of types of data and entities in a reference schema. In other words, the system can generate a model from a schema that includes an ontology defined based on relationships between entities and their attributes. The reference schema can be a system-provided or default reference schema, or alternatively a user-supplied or third-party reference schema.

[000397] Although some data integration frameworks may reverse engineer metadata from known source system types, they do not provide an analysis of the metadata to build a functional type system that can be used for pattern definitions and entity classifications. Harvesting metadata is also limited in scope and does not extend to profiling data for the extracted datasets or entities. The functionality to allow the user to specify a reference schema for ontology learning from which to build functional type system for use in complex process (business logic) and integration patterns in addition to entity classification (in similar topological space) is currently unavailable.

[000398] In accordance with an embodiment, one or more schemas can be stored in a reference HUB, which itself can be provided within or as part of the system HUB. As with the reference schema, a reference HUB can also be a user-supplied or third-party reference HUB, or in a multitenant environment, can be associated with a particular tenant, and accessed for example through a dataflow API.

[000399] In accordance with an embodiment, the reference HUB can be used to analyze data flows and to further classify or make recommendations for, e.g., transformations, enrichments, filtering, or cross-entity data fusion.

[000400] For example, in accordance with an embodiment, the system can receive input defining a reference HUB as a schema for ontology analysis. The reference HUB can be imported to obtain entity definitions (attributes definitions, data types, and relationships between datasets or entities, constraints, or business rules). Sample data (e.g., attribute vectors such as, for example, columnar data) in the reference HUB can be extracted for all datasets or entities, and profiled data, to derive several metrics on the data.

[000401] In accordance with an embodiment, a type system can be instantiated based on the nomenclature of the reference schema. The system can perform an ontology analysis, to derive an ontology (e.g., a set of rules) that describe types of data. The ontology analysis

can determine data rules that are defined terms of profiled data (e.g., attribute, or composite value) metrics and describe the nature of a business type element (e.g., UOM, ROI, or Currency Type) along with their data profile; relationship rules that define associations across datasets or entities, and attribute vectors (constraints or references imported from the reference schema); and complex rules that can be derived through a combination of data and relationship rules. A type system can then be defined based on the rules derived through metadata harvesting and data sampling.

5 [000402] In accordance with an embodiment, patterns and templates can be utilized from a system HUB, based on the type system instantiated using the ontology analysis. The system can then perform data flow processing using the type system.

10 [000403] For example, in accordance with an embodiment, classification and type annotation of datasets or entities can be identified by the type system of the registered HUB. The type system can be used to define rules for functional and business types derived from the reference schema. Using the type system, actions, such as, for example, blending, enriching, and transformation recommendations, can be performed on the entities identified in data flows based on the type system.

15 [000404] **Figure 44** illustrates a system for generating functional type rules, in accordance with an embodiment.

[000405] As illustrated in Figure 44, in accordance with an embodiment, a rule induction logic 850 or software component, which can be provided as software or program code that is executable by a computer system or other processing device, enables rules 851 to be associated with a functional type system 852

[000406] **Figure 44** illustrates a system for generating functional type rules, in accordance with an embodiment.

25 [000407] As illustrated in **Figure 45**, in accordance with an embodiment, a HUB 1 can act as a reference ontology, for use in type-tagging, comparing, classifying, or otherwise evaluating a metadata schema or ontology provided by other (e.g., newly-registered) HUBs, e.g., HUB 2 and HUB 3, and creating appropriate rules, for use by the Data AI system.

30 [000408] **Figure 46** illustrates an object diagram for use in generating functional type rules, in accordance with an embodiment.

[000409] In accordance with an embodiment, as illustrated in **Figure 46**, for example, the rule induction logic enables rules to be associated with a functional type system having a set of functional types 853 (for example, HUBs, datasets or entities, and attributes), and stored in registries for use in creating dataflow applications, e.g., pipelines, Lambda applications, including that each functional type 854 can be associated with functional type rules 856, and rules 858. Each rule can be associated with rule parameters 860.

[000410] In accordance with an embodiment, the first time a reference schema is processed, an ontology including a set of rules can be prepared that are appropriate to that schema.

[000411] In accordance with an embodiment, the next time a new HUB, or a new schema is evaluated, its datasets or entities can be compared with an existing ontology and prepared rules, and used in the analysis of the new HUB/schema and its entities, and the further learning of the system.

[000412] Although some metadata harvesting in data integration frameworks are limited to reverse engineering entity definitions (attributes and their data types and in some cases relationships); in accordance with an embodiment, the approach provided by the system described herein differs in allowing a schema definition to be used as reference ontology from which business and functional types can be derived, along with data profiling metrics for the datasets or entities, in the reference schema. This reference HUB can then be used to analyze business entities in other HUBs (data sources), to further classify or make recommendations (e.g., blending, or enriching).

[000413] In accordance with an embodiment, the system employs the following set of steps to ontology learning using a reference schema:

[000414] A user specifies the option to use a newly registered HUB as a reference schema.

[000415] Entity definitions (e.g., attributes definitions, data types, relationships between entities, constraints or business rules are imported).

[000416] A sample data is extracted for all datasets or entities, and data profiled to derive several metrics on the data.

[000417] A type system is instantiated (functional and business types), based on the nomenclature of the reference schema.

[000418] A set of rules that describe the business types are derived.

[000419] Data rules are defined in terms of profiled data metrics and describe the nature of a business type element (for example, UOM or ROI or Currency Type can be defined as a business type element along with their data profile).

[000420] Relationship rules are generated that define associations across elements (constraints or references imported from the reference schema).

[000421] Complex rules are generated that can be derived through a combination of data and relationship rules

[000422] The type system (functional and business) is defined based on the rules derived through metadata harvesting and data sampling.

[000423] Patterns or templates can then use the types instantiated based on the

reference schema to define complex business logic.

[000424] A HUB registered with the system can then be analyzed in the context of the reference schema.

[000425] Classification and type annotation of datasets or entities, in the newly registered HUB can be performed, based on the rules for functional and business types derived from the reference schema.

[000426] Blending, enriching, transformation recommendations can be performed on the datasets or entities, based on the type annotation.

[000427] **Figure 47** illustrates a process for generating a functional type system based on generated one or more rules, in accordance with an embodiment.

[000428] As illustrated in Figure 47, in accordance with an embodiment, at step 862, input is received defining a reference HUB.

[000429] At step 863, the reference HUB is accessed to obtain one or more entity definitions associated with datasets or entities, provided by the reference HUB.

[000430] At step 864, a sample data is generated for the one or more datasets or entities, from the reference HUB.

[000431] At step 865, the sample data is profiled to determine one or more metrics associated with the sample data.

[000432] At step 866, one or more rules are generated based on the entity definitions.

[000433] At step 867, a functional type system is generated based on the generated one or more rules.

[000434] At step 868, the functional type system, and profiles of the sample data, are persisted for use in processing a data input.

25 **Foreign Function Interface**

[000435] In accordance with an embodiment, the system provides a programmatic interface (referred to herein in some embodiments as a foreign function interface), by which a user or third-party can define a service, functional and business types, semantic actions, and patterns or predefined complex data flows based on functional and business types, in a declarative manner, to extend the functionality of the system.

[000436] As described previously, current data integration systems may provide limited interfaces with no support for types and no well-defined interfaces for object composition and pattern definition. Because of such shortcomings, complex functionality like cross service recommendations or unified application design platform to invoke semantic actions across services extending the framework is currently not available.

[000437] In accordance with an embodiment, the foreign function interface enables a

user to provide definitions or other information (e.g., from a customer other third-party) in a declarative manner, to extend the functionality of the system.

[000438] In accordance with an embodiment, the system is metadata-driven, such that definitions received through the foreign function interface can be processed to determine a metadata, determine a classification of the metadata, such as, for example, data types (e.g., functional and business types), and compare the data types (both functional and business) to existing metadata to determine whether there is a type match.

[000439] In accordance with an embodiment, the metadata received through the foreign function interface can be stored in a system HUB to be accessed by the system for processing a data flow. For example, the metadata can be accessed to determine semantic actions based on a type of data sets received as an input. The system can determine semantic actions permitted for the types of data provided through the interface.

[000440] In accordance with an embodiment, by providing a common declarative interface, the system can enable users to map service native types and actions to platform native types and actions. This allows a unified application design experience through type and pattern discovery. It also facilitates a purely declarative data flow definition and design involving components of various services extending the platform and generation of native code for the respective semantic actions.

[000441] In accordance with an embodiment, metadata received through the foreign function interface can be processed in an automated fashion so that the objects or artifacts described therein (e.g., data types or semantic actions) can be used in operation of data flows processed by the system. The metadata information received from the one or more third-party systems may also be used to define a service, indicate one or more functional and business types, indicate one or more semantic actions, or indicate one or more patterns/templates.

[000442] For example, in accordance with an embodiment, a classification of the accessed data can be determined, such as functional and business types of the data. The classification can be identified based on the information about the data received with the information. By receiving data from one or more third party systems, the functionality of the system can be extended to perform data integration on a data flow based on the information (e.g., a service, a semantic action, or a pattern) received from the third party.

[000443] In accordance with an embodiment, metadata in a system HUB can be updated to include the information identified about the data. For example, services and patterns/templates can be updated to be performed based on information (e.g., a semantic action) identified in the metadata received through the foreign function interface. Thus, the system can be augmented with functionality through the foreign function interface without interrupting processing of a data flow.

[000444] In accordance with an embodiment, subsequent data flows can be processed using the metadata in the system HUB after it is updated. Metadata analysis can be performed on a data flow of a dataflow application, e.g., pipeline, Lambda application. The system HUB can then be used to determine a recommendation of a transformation, taking into account the definitions provided via foreign function interface. The transformation can be determined based on a pattern/template, which is used to define semantic actions to perform for services, wherein the semantic actions can similarly take into account the definitions provided via foreign function interface.

[000445] **Figure 48** illustrates a system for identifying a pattern for use in providing a recommendation for a data flow, based on information provided via a foreign function interface, in accordance with an embodiment.

[000446] As illustrated in Figure 48, in accordance with an embodiment, definitions received via foreign function interface 900 can be used to update one or more of a service registry 902, functional and business type registry 904, or patterns/templates 906, within the system HUB.

[000447] In accordance with an embodiment, the updated information can be used by the data AI subsystem, including a rules engine 908, to determine, for example, type annotated HUBs, datasets or entities, or attributes 910 in the system HUB, and provide those datasets or entities, to a recommendation engine 912, for use in providing recommendations for a dataflow application, e.g., pipeline, Lambda application, via software development component (e.g., Lambda Studio).

[000448] **Figure 49** further illustrates identifying a pattern for use in providing a recommendation for a data flow, based on information provided via a foreign function interface, in accordance with an embodiment.

[000449] As illustrated in Figure 49, in accordance with an embodiment, a third-party metadata 920 can be received at the foreign function interface.

[000450] **Figure 50** further illustrates identifying a pattern for use in providing a recommendation for a data flow, based on information provided via a foreign function interface, in accordance with an embodiment.

[000451] As illustrated in Figure 50, in accordance with an embodiment, the third-party metadata received at the foreign function interface can be used to extend the functionality of the system.

[000452] In accordance with an embodiment, the system enables allows framework extensibility through well-defined interfaces that allow registration of services, types native to the services, semantic actions implemented by the services along with their typed parameters, patterns or templates that abstract predefined algorithms available as part of the service

among other things.

[000453] In accordance with an embodiment, by providing a common declarative programming paradigm, the pluggable service architecture allows the mapping of service native types and actions to platform native types and actions. This allows a unified application design experience through type and pattern discovery. It also facilitates a purely declarative data flow definition and design involving components of various services extending the platform and generation of native code for the respective semantic actions.

[000454] In accordance with an embodiment, the pluggable service architecture also defines a common interface for compilation, generation, deployment and runtime execution framework (Unified Application Design Service) for the plugins. The recommendation engine can machine learn and reason over semantic actions and patterns of all plugged in services and can make cross service semantic action recommendations for a distributed complex data flow design and development.

[000455] **Figure 51** illustrates a process for identifying a pattern for use in providing a recommendation for a data flow, based on information provided via a foreign function interface, in accordance with an embodiment.

[000456] As illustrated in Figure 51, in accordance with an embodiment, at step 932, one or more definitions of a metadata for use in processing a data are receive, via a foreign function interface.

[000457] At step 934, the metadata received via foreign function interface is processed, to identify information about the received metadata, including one or more of a classification, semantic action, template defining a pattern, or service defined by the received metadata.

[000458] At step 936, the metadata received via foreign function interface is stored in a system HUB, wherein the system HUB is updated to include the information about the received metadata and to extend functional capabilities of the system including its supported types, semantic actions, templates, and services.

[000459] At step 938, a pattern for providing a recommendation for a data flow is identified, based on the information updated in the system HUB via foreign function interface.

Policy Based Lifecycle Management

[000460] In accordance with an embodiment, the system can provide data governance functionality such as, for example, provenance (where a particular data came from), lineage (how the data was acquired/processed), security (who was responsible for the data), classification (what is the data about), impact (how impactful is the data to a business), retention (how long should the data live), and validity (whether the data should be excluded/included for analysis/processing), for each slice of data pertinent to a particular

snapshot in time; which can then be used in making lifecycle decisions and data flow recommendations.

[000461] Current approaches to managing data lifecycle does not involve tracking data evolution(changes in data profile or drift) or governance related functionality based on changes in data characteristics across temporal partitions. System observed or derived characteristics of data (classification, frequency of change, type of change, or use in processes) are not used to make lifecycle decisions or recommendations (retention, security, validity, acquisition intervals) on the data.

[000462] In accordance with an embodiment, the system can provide a graphical user interface that can indicate a lifecycle of data flow based on lineage tracking. The lifecycle can show where the data has been processed and if any errors have occurred during its processing; and can be shown as a timeline view of data (e.g., number of datasets, volume of datasets, and use of datasets). The interface can provide a point in time snapshot of data and can provide visual indicators for data as it is processed. As such, the interface can enable a complete audit of data, or a system snapshot of data based on the lifecycle (e.g., performance metrics, or resource usage).

[000463] In accordance with an embodiment, the system can determine a lifecycle of data based on sample data (periodically sampled from ingested data) and data acquired for processing by user defined applications. Some aspects of the data lifecycle management is similar across the category of ingested data namely streaming data and batch data (reference and incremental). For incremental data, the system can use scheduled, log collection, and event-driven methods to acquire temporal slices of data and manage allocation of slices across application instances covering the following functionality.

[000464] In accordance with an embodiment, the system can reconstruct data in case of loss using lineage across tiers from metadata managed in a system HUB.

[000465] For example, in accordance with an embodiment, incremental data attribute columns or user configured settings can be identified to acquire incremental data and maintaining high and low watermarks across ingesting data. Query or API and corresponding parameters (timestamp or Id columns) can be associated with data that is ingested.

[000466] In accordance with an embodiment, the system can maintain lineage information across tiers, such as, for example, query or log metadata in the edge layer, topic / partition offsets for each ingest in the scalable I/O layer, slices (file partitions) in the data lake, reference to process lineage (specific execution instance of the application producing the data and parameters associated with it) for subsequent downstream processed dataset using this data, topic/partition offset for the dataset "marked" to be published to a target endpoint and its corresponding data slice in the data lake, and publish job execution instance and the offsets

in the partitions that is processed and published to the target endpoint.

[000467] In accordance with an embodiment, in the case of failure of a layer (e.g., edge, scalable I/O, data lake, or publish), data can be reconstructed from the upstream layer or acquired from the source.

5 **[000468]** In accordance with an embodiment, the system can perform other lifecycle management functions.

[000469] For example, in accordance with an embodiment, security is enforced and audited at each of these layers for data slices. Data slices can be excluded or included (if already excluded) from being processed or accessed. This allows excluding spurious or
10 corrupt data slices from being processed. A retention policy can be enforced on slices of data through sliding windows. An impact is analyzed for slices of data (for example, the ability to tag slices for a given window, as being impactful in the context of a data mart built for quarterly reporting).

[000470] In accordance with an embodiment, data is classified by tagging the functional
15 or business types defined in the system (for example, tagging a dataset with a functional type (as a cube or dimensional or hierarchical data) along with the business type(s) (e.g., orders, customers, product, or time).

[000471] In accordance with an embodiment, the system can perform a method that includes accessing data from one or more HUBs. The data can be sampled, and the system
20 determine temporal slices of the data and managing the slices, including accessing a system HUB of the system, to obtain metadata about the sampled data. The sampled data can be managed for lineage tracking across one or more tiers in the system.

[000472] In accordance with an embodiment, incremental data and parameters about the sample data can be managed for the data that is ingested. The data can be classified by
25 tagging a type of data associated with the sample data.

[000473] **Figure 52** illustrates managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment.

[000474] For example, as illustrated in Figure 52, in accordance with an embodiment, the system can be used to receive a data from a HUB 952, in this example, an Oracle database,
30 and a HUB 954, in this example an S3 or other environment. Data received from input HUBs, at the edge layer, is provided to the scalable I/O layer, as one or more topics, for use by dataflow applications, e.g., pipelines, Lambda applications (wherein each of the topics can be provided as distributed partitions).

[000475] In accordance with an embodiment, the ingested data, typically represented by
35 an offset into a partition, can be normalized 964, by the compute layer, and written to the data lake as one or more temporal slices which span the tiers of the system.

[000476] In accordance with an embodiment, the data can then be used by downstream dataflow applications, e.g., pipelines, Lambda applications 966, 968, and ultimately published 970, to one or more additional topics 960, 962, and thereafter to a target endpoint (e.g., a table), at one or more output HUBs, such as in this example, a DBCS environment.

5 **[000477]** As illustrated in Figure 52, in accordance with an embodiment, at a first time, the data reconstruction and lineage tracking information can include information such as, for example, the provenance (Hub 1, S3), lineage (Source Entity in Hub 1), security (Connection Credential used), or other information regarding the ingest of data.

[000478] Figure 53 further illustrates managing a sampled data or accessed data, for
10 lineage tracking across one or more tiers, in accordance with an embodiment. As illustrated in Figure 53, at a subsequent time the data reconstruction and lineage tracking information can be updated to include information such as, for example, an updated provenance (→ T1), lineage (→ T1 (Ingest Process)), or other information.

[000479] Figure 54 further illustrates managing a sampled data or accessed data, for
15 lineage tracking across one or more tiers, in accordance with an embodiment. As illustrated in Figure 54, at a subsequent time the data reconstruction and lineage tracking information can be further updated to include information such as, for example, an updated provenance (→ E1), lineage (→ E1 (Normalize)), or other information.

[000480] In accordance with an embodiment, temporal slices 972, for use by one or more
20 dataflow applications, e.g., pipelines, Lambda applications, can be created, spanning the tiers of the system. In the event of a failure, for example a failure in writing to the data lake, the system can determine one or more unprocessed slices of data, and complete the processing of that slice of data, either in its totality, or incrementally.

[000481] Figure 55 further illustrates managing a sampled data or accessed data, for
25 lineage tracking across one or more tiers, in accordance with an embodiment. As illustrated in Figure 55, at a subsequent time the data reconstruction and lineage tracking information can be further updated, and additional temporal slices created, to include information such as, for example, an updated lineage (→ E11 (App1)), security (Role Executing App 1), or other information.

30 **[000482]** Figure 56 further illustrates managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment. As illustrated in Figure 56, at a subsequent time the data reconstruction and lineage tracking information can be further updated, and additional temporal slices created, to include information such as, for example, an updated lineage (→ E12 (App2)), security (Role Executing App 2), or other
35 information.

[000483] Figure 57 further illustrates managing a sampled data or accessed data, for

lineage tracking across one or more tiers, in accordance with an embodiment. As illustrated in Figure 57, at a subsequent time the data reconstruction and lineage tracking information can be further updated, and additional temporal slices created, to include information such as, for example, an updated lineage (→ T2 (Publish)), security (Role Executing Publish to I/O Layer), or other information.

[000484] Figure 58 further illustrates managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment. As illustrated in Figure 58, at a subsequent time the data reconstruction and lineage tracking information can be further updated, to reflect the output of the data to a target endpoint 976.

Data Lifecycle Management

[000485] In accordance with an embodiment, data lifecycle management based on the lineage tracking described above addresses several functional areas some of which can be configured by the user (access control, retention, validity), some derived (provenance, lineage) and others using machine learning algorithms (classification, impact). For example, the data management applies to both sample data (periodically sampled from ingested data) and data acquired for processing by user defined applications. Some aspects of the data lifecycle management is similar across the category of ingested data namely streaming data and batch data (reference and incremental). For incremental data, DFML uses scheduled, log collection, and event-driven methods to acquire temporal slices of data and manage allocation of slices across application instances covering the following functionality:

[000486] Reconstruction of data in case of loss using lineage across tiers from metadata managed in the system HUB.

[000487] Identifying of incremental data attribute columns or user configured settings to acquire incremental data and maintaining high and low watermarks across ingest.

[000488] Associating of query or API and corresponding parameters (timestamp or Id columns) for each ingest

[000489] Maintaining a lineage information across tiers. Query or Log metadata in the edge layer. Topic / Partition offsets for each ingest in the scalable I/O layer. Slices (file partitions) in the data lake. Reference to process lineage (specific execution instance of the application producing the data and parameters associated with it) for all subsequent downstream processed dataset using this data. Topic/Partition offset for the dataset "marked" to be published to a target endpoint and its corresponding data slice in the data lake. Publish job execution instance and the offsets in the partitions that is processed and published to the target endpoint.

[000490] In the case of failure of a layer, data can be reconstructed from the upstream

layer or acquired from the source. Security is enforced and audited at each of these layers for data slices. Data slices can be excluded or included (if already excluded) from being processed or accessed. This allows excluding spurious or corrupt data slices from being processed. Retention policy can be enforced on slices of data through sliding windows. Impact is analyzed for slices of data (for example, the ability to tag slices for a given window, as being impactful in the context of a data mart built for quarterly reporting).

[000491] Classification of the data by tagging the functional or business types defined in the system (for example, tagging a dataset with a functional type (as a cube or dimensional or hierarchical data) along with the business type(s) (e.g., orders, customers, product, or time).

[000492] Figure 59 illustrates a process for managing a sampled data or accessed data, for lineage tracking across one or more tiers, in accordance with an embodiment.

[000493] As illustrated in Figure 59, in accordance with an embodiment, at step 982, data is accessed from one or more HUBs.

[000494] At step 983, the accessed data is sampled.

[000495] At step 984, temporal slices are identified for sampled data or accessed data.

[000496] At step 985, a system HUB is accessed to obtain metadata about the sampled data or accessed data represented by the temporal slices.

[000497] At step 986, classification information is determined about the sampled data or accessed data represented by the temporal slices.

[000498] At step 987, the sampled data or accessed data represented by the temporal slices is managed, for lineage tracking across one or more tiers in the system.

[000499] Embodiments of the present invention can be implemented using one or more conventional general purpose or specialized digital computer, computing device, machine, or microprocessor, including one or more processors, memory and/or computer readable storage media programmed according to the teachings of the present disclosure. Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present disclosure, as will be apparent to those skilled in the software art.

[000500] In some embodiments, the present invention includes a computer program product which is a non-transitory computer readable storage medium (media) having instructions stored thereon/in which can be used to program a computer to perform any of the processes of the present invention. Examples of storage mediums can include, but are not limited to, floppy disks, optical discs, DVD, CD-ROMs, microdrive, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory devices, magnetic or optical cards, nanosystems (including molecular memory ICs), or other types of storage media or devices suitable for non-transitory storage of instructions and/or data.

[000501] The foregoing description of the present invention has been provided for the

purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations will be apparent to the practitioner skilled in the art.

5 **[000502]** For example while several of the embodiments described above illustrate the use of products such as, for example, Wolfram, Yago, Chronos, and Spark, to perform various computations, and data sources such as, for example, BDP, SFDC and S3, to act as sources or targets of data, the embodiments described herein can also be used with other types of products and data sources that provide similar types of functionalities.

10 **[000503]** Additionally, while several of the embodiments described above illustrate components, layers, objects, logic, or other or features of the various embodiments, such features can be provided as software or program code that is executable by a computer system or other processing device.

15 **[000504]** The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, thereby enabling others skilled in the art to understand the invention for various embodiments and with various modifications that are suited to the particular use contemplated. The modifications and variations include any relevant combination of the disclosed features. It is intended that the scope of the invention be defined by the following claims and their equivalents.

Claims:

What is claimed is:

- 5 1. A method for use with a data integration or other computing environment comprising:
accessing a data flow for each of one or more software applications;
processing the data flow for the one or more software applications to generate one or
more functional expressions representing the data flow, wherein the one or more functional
expressions are generated based on one or more semantic actions or business rules identified
10 in the data flow;
identifying a pattern of transformation in the data flow, for the one or more functional
expressions generated for each of the one or more software applications, wherein the one or
more semantic actions or business rules are used to identify the pattern of transformation in
the data flow; and
15 using the pattern of transformation identified in the data flow, providing a
recommendation of one or more data transformations for a data flow of another software
application.
- 20 2. The method of Claim 1, wherein the application is a dataflow application.
3. The method of Claim 1 or 2, wherein the pattern is used in displaying, at a graphical
user interface, selected ones of the semantic actions enabled for the accessed data, for
selection and use with the accessed data, including automatically providing or updating a list
of the selected ones of the semantic actions enabled for the accessed data, during the
25 processing of the accessed data.
4. The method of any of Claims 1 to 3, wherein an application represents a top level data
flow transformation; an action represents an operator on one or more datasets; and actions
reference a base semantic action or function declaratively defined in the system, and can have
30 one or more action parameters.
5. The method of Claim 4, wherein action parameters are owned by actions and have
specific functional or business types, and represent a specific upstream dataset to be
processed.
35
6. The method of Claim 4, wherein a scope resolver is used to resolve the reference to a

particular attribute or embedded object.

7. The method of any of Claims 1 to 6, wherein the method is performed in a cloud or cloud-based computing environment.

5

8. A system for providing recommendations of actions and transformations on an input data for use with a data integration or other computing environment, comprising:

one or more processors operable to:

access a data flow for each of one or more software applications;

10

process the data flow for the one or more software applications to generate one or more functional expressions representing the data flow, wherein the one or more functional expressions are generated based on one or more semantic actions or business rules identified in the data flow;

15

identify a pattern of transformation in the data flow, for the one or more functional expressions generated for each of the one or more software applications, wherein the one or more semantic actions or business rules are used to identify the pattern of transformation in the data flow; and

20

use the pattern of transformation identified in the data flow, providing a recommendation of one or more data transformations for a data flow of another software application.

9. The system of Claim 8, wherein the application is a dataflow application.

25

10. The system of Claim 8 or 9, wherein the pattern is used in displaying, at a graphical user interface, selected ones of the semantic actions enabled for the accessed data, for selection and use with the accessed data, including automatically providing or updating a list of the selected ones of the semantic actions enabled for the accessed data, during the processing of the accessed data.

30

11. The system of any of Claims 8 to 10, wherein an application represents a top level data flow transformation; an action represents an operator on one or more datasets; and actions reference a base semantic action or function declaratively defined in the system, and can have one or more action parameters.

35

12. The system of Claim 11, wherein action parameters are owned by actions and have specific functional or business types, and represent a specific upstream dataset to be

processed.

13. The system of Claim 11, wherein a scope resolver is used to resolve the reference to a particular attribute or embedded object.

5

14. The system of any of Claims 8 to 13, wherein the system is provided in a cloud or cloud-based computing environment.

15. A non-transitory computer readable storage medium, including instructions stored thereon which when read and executed by one or more computers cause the one or more computers to perform a method comprising:

10

accessing a data flow for each of one or more software applications;

processing the data flow for the one or more software applications to generate one or more functional expressions representing the data flow, wherein the one or more functional expressions are generated based on one or more semantic actions or business rules identified in the data flow;

15

identifying a pattern of transformation in the data flow, for the one or more functional expressions generated for each of the one or more software applications, wherein the one or more semantic actions or business rules are used to identify the pattern of transformation in the data flow; and

20

using the pattern of transformation identified in the data flow, providing a recommendation of one or more data transformations for a data flow of another software application.

25

16. The non-transitory computer readable storage medium of Claim 15, wherein the application is a dataflow application.

30

17. The non-transitory computer readable storage medium of Claim 15 or 16, wherein the pattern is used in displaying, at a graphical user interface, selected ones of the semantic actions enabled for the accessed data, for selection and use with the accessed data, including automatically providing or updating a list of the selected ones of the semantic actions enabled for the accessed data, during the processing of the accessed data.

35

18. The non-transitory computer readable storage medium of any of Claims 15 to 17, wherein an application represents a top level data flow transformation; an action represents an operator on one or more datasets; and actions reference a base semantic action or function

declaratively defined in the system, and can have one or more action parameters.

19. The non-transitory computer readable storage medium of Claim 18, wherein action parameters are owned by actions and have specific functional or business types, and represent
5 a specific upstream dataset to be processed.
20. The non-transitory computer readable storage medium of Claim 18, wherein a scope resolver is used to resolve the reference to a particular attribute or embedded object.
- 10 21. A computer program comprising program instructions in machine-readable format that when executed by a computer system cause the computer system to perform the method of any of Claims 1 to 7.
22. A computer program product comprising the computer program of Claim 21 stored in a
15 non-transitory machine readable data storage medium.
23. An apparatus comprising means for performing the method of any of Claims 1 to 7.




FIGURE 1




FIGURE 2




FIGURE 3




FIGURE 4




FIGURE 5




FIGURE 6




FIGURE 7




FIGURE 8




FIGURE 9




FIGURE 10




FIGURE 11




FIGURE 12




FIGURE 13




FIGURE 14




FIGURE 15




FIGURE 16




FIGURE 17




FIGURE 18




FIGURE 19




FIGURE 20




FIGURE 21




FIGURE 22




FIGURE 23




FIGURE 24




FIGURE 25




FIGURE 26




FIGURE 27




FIGURE 28




FIGURE 29




FIGURE 30




FIGURE 31




FIGURE 32




FIGURE 33




FIGURE 34





FIGURE 35



Display, at a graphical user interface, selected ones of the semantic actions enabled for the accessed data, for selection and use with the accessed data, including automatically providing or updating a list of the selected ones of the semantic actions enabled for the accessed data, during the processing of the accessed data.

FIGURE 36




FIGURE 37




FIGURE 38




FIGURE 39




FIGURE 40




FIGURE 41




FIGURE 42




FIGURE 43




FIGURE 44




FIGURE 45

Data AI Subsystem 168

Rule Induction Logic 850




FIGURE 46




FIGURE 47




FIGURE 48




FIGURE 49




FIGURE 50




FIGURE 51




FIGURE 52




FIGURE 53




FIGURE 54




FIGURE 55




FIGURE 56




FIGURE 57




FIGURE 58




FIGURE 59

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/048044

A. CLASSIFICATION OF SUBJECT MATTER
 INV. G06N5/02
 ADD. G06F9/54 G06F9/50

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 G06N G06F

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
 EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2016/049460 A1 (ORACLE INT CORP [US]) 31 March 2016 (2016-03-31) abstract; claims 1,8; figures 1,2,3,4a,5a,6	1-23
X	----- Michael Malak (Oracle): "Extending Word2Vec for Performance and Semi Supervised Learning", youtube, 24 June 2015 (2015-06-24), page 1 pp., XP054977814, Retrieved from the Internet: URL:https://www.youtube.com/watch?v=A2XNln XGeSY [retrieved on 2017-10-23] Slides and Audio passages at 00:41-00:59; 01:00-01:33; 2:50-3:20; 8:30-9:44; 12:00-14:29; 17:15-17:49 ----- -/--	1-23

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search 31 October 2017	Date of mailing of the international search report 07/11/2017
--	--

Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Manfrin, Max
--	--

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/048044

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	<p>Anonymous: "Apache Spark 2.0 Preview: Machine Learning Model Persistence - The Databricks Blog",</p> <p>³ 31 May 2016 (2016-05-31), XP055417721, Retrieved from the Internet: URL:https://databricks.com/blog/2016/05/31/apache-spark-2-0-preview-machine-learning-model-persistence.html [retrieved on 2017-10-20] the whole document</p> <p>-----</p>	1-23
Y	<p>Anonymous: "ML Pipelines: A New High-Level API for MLlib - The Databricks Blog",</p> <p>⁷ 7 January 2015 (2015-01-07), XP055417730, Retrieved from the Internet: URL:https://databricks.com/blog/2015/01/07/ml-pipelines-a-new-high-level-api-for-ml-lib.html [retrieved on 2017-10-20] the whole document</p> <p>-----</p>	1-23
A	<p>WO 2004/021186 A2 (SAP AG [DE]) 11 March 2004 (2004-03-11) abstract; figures 10-16</p> <p>-----</p>	1,8,15, 21,23

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2017/048044

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2016049460	A1	31-03-2016	
		CN 106796595 A	31-05-2017
		EP 3198484 A1	02-08-2017
		US 2016092474 A1	31-03-2016
		WO 2016049460 A1	31-03-2016

WO 2004021186	A2	11-03-2004	
		AU 2003260462 A1	19-03-2004
		WO 2004021186 A2	11-03-2004
