



(51) International Patent Classification:
G06F 16/26 (2019.01)

(21) International Application Number:
PCT/US2018/062154

(22) International Filing Date:
21 November 2018 (21.11.2018)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/594,938 05 December 2017 (05.12.2017) US
15/962,374 25 April 2018 (25.04.2018) US

(71) Applicant: **MICROSOFT TECHNOLOGY LICENSING, LLC** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventors: **SOBHY DERAZ, Ehab**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **EL-**

LIS, Charles David; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **WOOLF, Michael Joe**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **ASPLUND, Matthew Willem**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **OLINGER, Mathew B.**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **JAGDALE, Sameer Surendra**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **STREMPER, Charles**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(74) Agent: **MINHAS, Sandip S.** et al.; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(54) Title: MODULAR DATA INSIGHT HANDLING FOR USER APPLICATION DATA




FIG. 1

(57) **Abstract:** In non-limiting examples of the present disclosure, systems, methods and devices for assisting with providing dataset insights associated with datasets are presented. A dataset and a query relating to the dataset may be received. The dataset may be processed to determine metadata that describes one or more properties of the dataset. The dataset, the determined metadata, and the user query may be provided to one or more modular recommendation elements for processing into an insight result that indicates a result from data analysis directed to the query. The insight result may be transferred in a portable format for use by the productivity application in displaying one or more insight objects based on the insight result.

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, 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).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

MODULAR DATA INSIGHT HANDLING FOR USER APPLICATION DATA

BACKGROUND

[001] Various user productivity applications allow for data entry and analysis. These applications can provide for data creation, editing, and analysis using spreadsheets, presentations, documents, messaging, or other user activities. Users can store data files associated with usage of these productivity applications on various distributed or cloud storage systems so that the data files can be accessible wherever a suitable network connection is available. In this way, a flexible and portable user productivity application suite can be provided.

[002] However, the information technology industry has continually increased the amount of information as well as the quantity of sources of information. Users can be quickly overwhelmed with data analysis due to the sheer quantity of data or number of options available for managing and presenting the data and associated analysis conclusions. Moreover, users within an organization have a difficult time leveraging the data and analysis of co-workers, and leveraging data analysis while switching between small form-factor devices (such as smartphones and tablet computers) and large form-factor devices (such as desktop computers).

SUMMARY

[003] This summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter. Additional aspects, features, and/or advantages of examples will be set forth in part in the description which follows and, in part, will be apparent from the description or may be learned by practice of the disclosure.

[004] Non-limiting examples of the present disclosure describe systems, methods and devices for providing dataset insights for a productivity application. A query and/or command related to a dataset may be received. The dataset may comprise a plurality of values included in the productivity application and/or a plurality of values from one or more secondary data sources. The query and/or command may comprise a user-provided input to have an insight related to the dataset surfaced on a display of a user interface of the productivity application or an automated query and/or command generated by the productivity application (e.g., queries and/or commands may be automatically generated

when new datasets are created).

[005] In examples, the query and/or command, the dataset and metadata associated with the dataset may be sent to an insight platform for processing. The insight platform may identify, based on the received query and/or command, dataset, and metadata, one more insight types for surfacing via the productivity application. In some examples, the insight platform may identify insight types for surfacing based on additional information, including user preferences, the dataset type, a type of information included in the dataset, past user interactions with the insight platform, past user interactions with insights, past user interactions with the productivity application, and/or past user interactions with one or more datasets related to the received dataset.

[006] In some examples, the insight platform may send one or more insight types to the productivity application in a portable format for surfacing on the display of the productivity application. The surfaced one or more insights may comprise graphs, charts, tables, pivot tables, textual content, audio content, and/or pictures. The information included in the one or more surfaced insights may comprise visual and/or audio representations of: information from the dataset, summaries of information from the dataset, comparison of information from the dataset with information from one or more other datasets, projections associated with information from the dataset, and selectable elements for providing information about the surfaced insights based on receiving a modification to the one or more insights, receiving a modification to one or more values included in a dataset related to the one or more insights, and/or receiving a query and/or command directed to the generation of the one or more insights.

BRIEF DESCRIPTION OF THE DRAWINGS

[007] Many aspects of the disclosure can be better understood with reference to the following drawings. While several implementations are described in connection with these drawings, the disclosure is not limited to the implementations disclosed herein. On the contrary, the intent is to cover all alternatives, modifications, and equivalents.

[008] Figure 1 illustrates a data insight environment in an example.

[009] Figure 2 illustrates operations of data insight environments in an example.

[0010] Figure 3 illustrates operations of data insight environments in an example.

[0011] Figure 4 is a first exemplary method for providing insight results in a productivity application.

[0012] Figure 5 is a second exemplary method for providing dataset insights for a productivity application.

[0013] Figure 6 illustrates a computing system suitable for implementing any of the architectures, processes, and operational scenarios disclosed herein.

DETAILED DESCRIPTION

[0014] User productivity applications provide for user data creation, editing, and analysis using spreadsheets, slides, documents, messaging, or other application activities.

However, due in part to continually increasing amounts of user data as well as the quantity of different sources of information, users can be quickly overwhelmed with tasks related to analyzing this data. In workplace environments, such as a company or other organization, users might have a difficult time leveraging the data and analysis performed by other co-workers. This level of growth in data analysis increases a need to augment a user ability to make sense and use increasing sources and volumes of data.

[0015] In the examples herein, user data can be leveraged in various data visualization environments to create "insight" results or recommendations for users during data analysis stages. In some examples, insight results, as described herein, may comprise extensions of analytic objects that include charts, pivot tables, tables, graphs, and the like. In additional examples, insight results may comprise further content that represents an insight, such as summary verbiage, paragraphs, graphs, charts, pivot tables, data tables, or pictures that are generated for users which indicate key takeaways from the data.

[0016] Turning now to a first example system for data visualization and insight generation, Figure 1 is presented. Figure 1 illustrates data visualization environment 100. Environment 100 includes user platforms 110 and insight platform 120. Each of the elements of environment 100 can communicate over one or more communication links, which can comprise physical and logical network links, among others.

[0017] User platform 110 provides user interface 112 to application 111. Application 111 can comprise a user productivity application for use by an end user in data creation, analysis, and presentation. For example, application 111 might comprise a spreadsheet, word processor, database, or presentation application. User platform 110 also includes insight module 114. Insight module 114 can interface with insight platform 120 as well as provide insight services within application 111. User interface 112 can include graphical user interfaces, console interfaces, web interfaces, text interfaces, among others.

[0018] Insight platform 120 provides insight services, such as insight service 121, insight application programming interface (API) 122, metadata handler 123, and recommendation platform 124. Insight service 121 can invoke various other elements of insight platform 120, such as API 122 for interfacing with clients. Insight service 121 can

also invoke one or more recommendation modules, such as provided by recommendation platform 124.

[0019] In operation, insight service 121 in coordination with API 122, metadata handler 123, and recommendation platform 124 can process one or more datasets to establish data insight results, referred to in Figure 1 as portable insights 144, that can be provided to clients/user platforms which present graphical visualization portions, data descriptions or conclusions/summaries, object metadata, as well as the underlying datasets. Portable insights 144 can produce extensions of typical analytic objects, such as charts, graphs, tables, pivot tables, data descriptions, and other data or document presentation elements. Portable insights 144 can include other content that represents insight objects, such as verbiage or summary statements that provide additional information to a user, such key takeaways of data insight analysis, and other data descriptions.

[0020] In operation, a user of user platform 110 or application 111 might indicate a set of data or target dataset for which data insight analysis is desired. This analysis can include traditional data analysis such as math functions, static graphing of data, pivoting within pivot tables, or other analysis. However, in the examples herein, an enhanced form of data analysis is performed, namely insight analysis. At the user application level, one or more insight modules are included to not only present insight analysis options to the user but also interface with insight platform 120 which performs the insight analysis, among other functions. Upon designation of one or more target datasets, a user can employ insight service 121 via insight API 122 to process the target datasets and generate one or more candidate insights, portable insight results, and associated insight metadata. In Figure 1, this process is shown using user data 141 and optional metadata 142 supplied by the user and/or application 111. However, it should be understood that target datasets can be supplied from other data sources, including in-application data sources, data documents, data storage elements, distributed data storage systems, or other data sources, such as data repositories.

[0021] As mentioned above, metadata 142 can also be provided with the user data 141. This metadata might be omitted in some examples, and determined by metadata handler 123 of insight platform 120. Metadata 142 can include properties or descriptions about user data 141, such as column/row headers, data contexts, application properties, and other information. Moreover, identifiers can be associated with the user data or with already-transferred user data and metadata. These identifiers can be used by insight module 114 to reference the data/metadata within insight platform 120. A further discussion of these

identifiers is discussed below. Metadata processing of metadata handler is discussed in Figures 2-3 below.

[0022] Metadata handler 123 processes user data sets, such as user data 141, along with any user-provided or application-provided metadata associated with user data 141.

5 Metadata handler 123 determines various metadata associated with user data 141, such as extracting properties, data descriptions, headers, footers, column/row descriptors, or other information. For example, if data 141 is provided as a table with column and/or row headers, metadata handler 123 can extract the column or row headers as metadata.

Moreover, metadata handler 123 can intelligently determine what the column/row
10 information metadata might comprise in examples where metadata accompanies the data or when metadata does not accompany the data. For example, properties of the data itself might be determined by metadata handler 123 to establish metadata for the data, such as data features, numerical formats, symbols embedded with the data, patterns among the data, column or row organizations determined for the data, or other data properties.

15 Metadata 142 that might accompany user data 141 can also inform further metadata analysis by metadata handler 123, such as when only a subset of the data is labeled or has headers.

[0023] Once metadata is determined for the data sets, Metadata handler 123 can cache or otherwise store the metadata, along with any associated user data, in cache 132. Cache

20 132 can comprise one or more data structures for holding metadata and user data for use by insight service 121 and recommendation platform 124. Cache 132 can advantageously hold this data/metadata for use over one or more insight analysis processes and user application requests for analysis. Various identifiers can be associated with the data or metadata for reference by insight module 114 when performing further/later data insight

25 analysis. Insight results determined for various user data sets can also be stored in association with the identifiers for later retrieval, referencing, or handling by any module, service, or platform in Figure 1. Moreover, metadata and data cached in cache 132 can be employed in parallel by any of recommendation modules 130. In some examples one or more components of insight platform 120 (e.g., insight service 121, metadata handler 123,
30 recommendation platform 124, insight API 122) may send metadata 142 back to user platform 110 upon metadata handler 123 determining properties associated with metadata 142. Metadata 142, and properties associated with metadata 142, may be stored in association with application 111 and/or a document containing a data set from which metadata 142 was determined. Thus, in examples where metadata 142 is sent back to user

platform 110, insight module 114 may not have to communicate with cache 132 for further/late data insight analysis of a previously analyzed data set because metadata 142 is stored with application 111 in user platform 110.

[0024] Insight service 121 establishes content of the data insight results according to processing a target user dataset using data analysis recommenders provided by recommendation platform 124. Portable insights 144 can indicate insight results and insight candidates for presentation to a user by application 111. For example, portable insights 144 can describe insight results in a manner that can be interpreted by application 111 to produce application-specific insight objects for presentation to a user. These insight objects can be presented in user interface 112, such as for inclusion into a spreadsheet canvas of a spreadsheet application. Object metadata, such as metadata determined by metadata handler 123, can accompany portable insights 144.

[0025] To determine the data insight results, recommendation modules 130 are employed. These recommendation modules can be used to establish data analysis preferences derived from past user activity, application usage modalities, organizational traditions with regard to data analysis, individualized data processing techniques, or other activity signals.

Knowledge graphing or graph analysis can be employed to identify key processes or data analysis techniques that can be employed in the associated insight analysis. Knowledge repositories can be established to store these data analysis preferences and organizational knowledge for later use by users employing the insight services discussed herein.

Machine learning, heuristic analysis, or other intelligent data analysis services can comprise the recommendation modules. Each module can be 'plugged' into recommendation platform 124 for use in data analysis to produce insight recommendations for the user data. Recommendation modules 131-133, among others, might be dynamically added or removed, instantiated or de-instantiated, among other actions, responsive to the user data, metadata, desired analysis types, user instructions, application types, past analyses on user data, or other factors.

[0026] Turning now to a further discussion of recommendation platform 124, insight service 121 can grow to support one or more recommenders and recommendation types. It can be advantageous to call out the explicit integration steps for recommenders to hook into insight service 121. Below are example processes by which a new recommender may register itself, as well as a processing pipeline for creating machine-learned intelligent recommendation providers.

[0027] Several terms are included in the discussion herein, which have example

descriptions as follows. Featurization [Feature Extraction] is a machine learning term used to describe the process of converting the raw input into a collection of features used as inputs into a Machine Learning Model. A 'feature' comprises an individual measurement used as input into a Machine Learning Model. Metadata can include information describing general properties about a given dataset, such as column types, data orientation, and the like. Lazy Evaluation comprises a process by which a value is only calculated when explicitly requested. A recommender may comprise a single algorithm, either heuristic or machine-learning based, that takes in the provided metadata from a dataset, and generates a set of recommendations, such as charts, tables, design, and the like. Through the application of featurization and machine learning, recommenders can be intelligently trained to identify data structures and/or metadata associated with datasets that the recommenders can generate insights for in association with insight platform 120. Featurization and machine learning may be applied on an entity-specific basis, such that insight types (e.g., charts, tables, design) that entities (e.g., individual users, user demographics, corporate entities, entity groups) have indicated a preference for over time may be generated by appropriate recommenders. Thus, through the training of recommenders, and the application of lazy evaluation, only values that are associated with recommenders that generate insight types that are relevant/preferred to specific entities need be calculated, thereby significantly reducing the processing costs associated with calculation of values related to non-preferred recommenders, and storage costs associated with caching or otherwise storing values for recommenders that are not relevant to the entities.

[0028] During usage of recommendation platform 124, sharing allows for easy sharing of as much code and resources as possible between training, testing and production. This can be achieved using shared binaries and shared processing pipelines. Versioning allows for easily changing the versions of all parts of the pipeline and ensuring they can be kept in sync. Quality controls maintain a minimum quality bar for all recommendation modules for both accuracy and performance.

[0029] The development of a recommendation module can be broken down into three main stages: generation, validation and production. The generation stage consists of either training a machine learning model or designing/implementing a heuristic-based algorithm. Once a recommendation module is created it can be run through multiple rings of validation. This may consist of both a performance portion and a quality assurance. Each recommendation module can be assigned a budget for processor time as well as minimum

required accuracy. The final stage of the pipeline can include running each individual recommendation module in production. Recommendation platform 124 can be responsible for federating out individual requests to all registered recommendation modules and aggregating the results.

5 **[0030]** This design advantageously supports the ability for machine learning models to be trained on a feature set that is as identical as reasonable to what may be seen in real user data. This means that as updates are made to the supported recommendation module feature set and associated generation logic, each recommendation module can train a new model that can be utilized to match the new version and the production service can ensure
10 that the hosted models are in sync with their feature set version. A part of recommendation platform 124 is the continued improvement and expansion of the features. To ensure that the machine learning/training models are working as expected, the same logic may be used to generate the features that are used to train the models as well as validate and run them.

15 **[0031]** Turning now to the operation of API 122, various inputs and outputs are provided. As input, API 122 can receive user data 141, such as datasets in a two-dimensional tabular format. In some examples, this user data may have accompanying metadata 142. In other examples this user data may have embedded metadata. In still other examples, this user data may have no accompanying metadata. One or more applications and/or users
20 associated with the infrastructure described herein may initiate one or more queries or questions posed towards user data 141. These queries are indicated in queries 143 in Figure 1, and can comprise natural language questions posed by users and/or applications related to the user data and submitted through API 122 in a standardized format. A user might ask one or more questions for analysis by insight platform 120, and provide a
25 portion of data to insight platform 120. The queries indicated by the user and/or application, and included in queries 143, can include questions such as, "I need charts for this data..." or "Provide the metadata for this data..." or "Summarize this data..." among other query types. API 122 provides for input mechanisms for application 111 through insight module 114 to input the data, metadata, and queries for use by insight service 121.
30 Based on the inputs (e.g., data, metadata, queries), insight platform 120 provides, through API 122, one or more insight results indicated by portable insights 144.

[0032] As outputs, such as portable insights 144, API 122 can provide insight results in a standardized output for interpretation by any application to present the insight results to the user in that application's native format. Portable insights 144 comprises descriptions

of the insight results which can be interpreted by application 111 or insight module 114 to generate visualizations of the insight results to users. In this manner, a flexible and/or portable insight result can be presented as an output by API 122 and interpreted for display as-needed and according to specifics of the application user canvas.

5 **[0033]** API 122 defines the formatting for inputs and outputs, so that applications and users can consistently present data, metadata, and queries for analysis by insight platform 120. API 122 also defines the mechanisms by which application 111 can communicate with insight platform 120, such as allowed input types, input ranges, and input formats, as well as possible outputs resultant from the inputs. API 122 also can provide identifiers
10 responsive to provided user data, metadata, and queries so that data, metadata, and queries can be referenced later by clients, such as application 111, as stored in cache 132.

[0034] In one example, API 122 comprises an insights representational state transfer (REST) style of API. The insights REST API comprises a web service for applying heuristic and machine learning-based analysis to a set of data in order to retrieve high level
15 interesting views, called insights herein, of the data. The insights REST API can provide recommendations for charts and/or pivots of the user data. The insights REST API can also provide metadata services used for natural language insights, and other analysis.

[0035] An example operation flow involving a client, such as application 111, communicating with API 122 may comprise the following flow.

20 **[0036]** At a first operation, a client uploads a range of client data to the service, which initiates a data session. In some examples, this may cause a URL to be returned containing a unique “range id” that is 1:1 with the data session. In examples where a user triggered refresh has occurred, a new “range id” may be generated and returned in a URL.

[0037] At a second operation, the client provides an indication of a type of analysis they
25 want performed. Analysis options may include receiving recommendations for insights, or metadata services used for natural language insights, among other analysis choices. This returns an Operation ID, which is 1:1 with the process of performing the requested analysis.

[0038] At a third operation, the client waits for the operation to complete, periodically
30 polling the service, and at a fourth operation the client is provided with an opportunity to cancel an operation.

[0039] At a fifth operation, the client gets the results of the completed operation. Additional requests may be made on the same data in cache 132 (e.g., correct the metadata as requested by the user, then get new recommendations), without needing to upload the

data again. That is, the operation flow may return to the second operation.

[0040] At a sixth operation, the client closes the data session, and the data session ends.

[0041] To illustrate example data set handling and metadata determination, Figure 2 is presented. Figure 2 illustrates further operations of the elements of Figure 1, although the

5 operations of Figure 2 can be implemented by elements other than those of Figure 1. In operation, dataset 200 can be provided along with one or more queries 201 directed to the dataset to an insight platform. For example, dataset 200 and query 201 might be provided via API 122 for processing by insight platform 120 of Figure 1. Insight platform 120 can process the dataset and query to provide an insight result, which can be interpreted by
10 application 111 for display as insight objects 202.

[0042] In Figure 2, example dataset 200 is shown comprising a two-dimensional array of data in a spreadsheet application user interface. In examples, dataset 200 can comprise a table, pivot table, spreadsheet, or other dataset, or can be a subset thereof. As seen in Figure 2, dataset 200 comprises data along with metadata. The data comprises user data

15 values or user data entered for analysis. The metadata includes descriptions of the data, which in this case is column headers which indicate properties of the data contained in underlying columns. The metadata in dataset 200 indicates a first column "name" and a second column "score." When submitted through API 122, insight service 121 can employ metadata handler 123 to isolate the metadata from the data, along with

20 determining other metadata as appropriate. The data and the metadata can be stored in association with an identifier into cache 132. Metadata handler 123 can provide table detection services for provided datasets. These table detection services can detect not only data arranged into two-dimensional arrays, such as tables, but also extract metadata that describes the data in the arrays.

25 **[0043]** Insight service 121 can initiate insight processing for the dataset using the metadata and one or more recommendation modules (e.g., recommendation modules 131-133). These recommendation modules can process the datasets, the queries, and the metadata to determine one or more insight results using machine learning techniques, heuristic processing, natural language processing, artificial intelligence services, or other
30 processing elements. The insight results, as discussed herein, are presented in a portable description format, such as using a markup language (e.g. HTML, XML). A user application comprising insight handling functions can interpret the insight results in the portable format and generate one or more insight objects for rendering into a user interface and presentation to a user.

[0044] An exemplary portable insight client/application interaction, utilizing insight service 121 and API 122, is described below:

- Insight module 114 sends data to insight service 121. Insight service 121 replies with a location for RESTful resource tracking of the data.
- 5 • Insight module 114 tells insight service 121 to generate insight recommendation results and that the application is capable of rendering charts and PivotCharts. A long running task will be created on insight service 121 and insight service 121 replies with a RESTful resource that insight module 114 can use to track this operation.
- 10 • Insight module 114 queries state of operation and is told that the operation is running. Insight module 114 is also told to try polling again after a specified time lapse.
- Insight module 114 queries state of operation later and is told that the operation has succeeded. Insight module 114 is also given the location of
15 the created resource.
- Insight module 114 asks for the insight recommendation results. In this example, there are two PivotChart recommendations, notably insight results that correspond to insight objects 202.
- Insight module 114 tells insight service 121 that insight module 114 is done
20 with the resource tracking the data. In some examples, insight service 121 may store this data for a short amount of time (on the order of hours). In other examples, the notification that insight module 114 is done with the resource tracking of the data provides insight service 121 a request to clean up the resource immediately, thereby increasing storage capacity of one or
25 more devices where the resource tracking data is stored.

[0045] As a further example involving the elements of Figure 1, application 111 can comprise a spreadsheet application, word processing application, presentation application, or other user application. Application 111 may comprise various user interface elements presented by user interface 112, such as windowed dialog boxes, a user canvas from which
30 data can be entered and manipulated, various menus, icons, control elements, and/or status informational elements. Furthermore, insight module 114 provides for enhanced user interface elements from which a user can initiate insight processing by insight platform 120, such as responsive to a user selecting an insight trigger icon or entering an insight

analysis command. In some examples, users may provide background services with authorization to monitor target data sets, which can be utilized to pre-compute insight results for presentation to a user.

[0046] Typically, a user may have a set of data entered into a worksheet or other

5 workspace presented by application 111. This data can comprise one or more structured tables of data, and/or unstructured data, and can be entered by a user or imported from other data sources into the workspace. A user might want to perform data analysis on this target data, and can select among various data analysis options presented by user interface 112. However, typical options presented for data analysis by user interface 112 and the
10 associated application 111 might only include static graphs or might only include content that the user has manually entered. This manual content can include graph titles, graph axes, graph scaling, colors, and/or other graphical and textual content or formatting.

[0047] Example insight generation operations proceed according to a modular analysis provided by recommendation modules 130. Insight service 121 can instantiate, apply, or

15 otherwise employ one of recommendation modules 130 to perform the insight analysis. As discussed herein, this can include analysis processes that are derived by processing metadata, query structure and content, along with other data which might comprise past usage activities, activity signals, and/or usage modalities that are found in the data. The target dataset can be processed according to various formulae, equations, functions, and
20 the like to determine patterns, outliers, majorities/minorities, segmentations, and/or other properties of the target dataset that can be used to visualize the data and/or present conclusions related to the target dataset. Many different analysis processes can be performed in parallel.

[0048] Insight results are determined by recommendation modules 130 and provided to

25 insight service 121 for various formatting and standardization into the portable format output by API 122. API 122 can provide these portable insights for delivery to insight module 114 of application 111. Insight module 114 can interpret the insight results in the portable format to customize, render, or otherwise present the insight results to a user in the application. For example, when the insight results procedurally describe charts,
30 graphs, or other graphical representations of insight results, the application can present these graphical representations.

[0049] In Figure 2, insight results can be rendered into insight objects 202, such as the two charts shown. Metadata extracted or determined for the dataset can be included in the insight results/objects to label axes, label data portions, or otherwise provide context and

descriptions for the insight results/objects. The selection or choice of an object type, such as graph or chart type, can be determined based on the dataset content, the metadata, or according to the query presented, among other considerations. For example, the query might indicate that a graph or chart, or particular graph/chart types are to be provided.

5 [0050] These can be presented in a graphical list format, paged format, or other display formats that can include further insights objects available via scrollable user interface operations or paged user interface operations. A user can select a desired insight object, such as a graph object, for insertion into a spreadsheet or other document. Once inserted, further options can be presented to the user, such as dialog elements from which further
10 insights can be selected. Each insight object can have automatically determined object types, graph types, data ranges, summary verbiage, supporting verbiage, titles, axes, scaling factors, or color selections, or other features. These features can be determined by the recommendation modules using the insight results discussed herein, such as based on data analysis derived from data, metadata, or queries.

15 [0051] Further options can be presented to the user that allow for secondary manipulation of the insight objects or insight results. Secondary manipulation can include manipulation of the dataset or metadata to perform further insight analysis. Secondary manipulation can include various queries or questions that a user can ask about the insight object presently presented to the user, such as questions including "what happened," "why did this
20 happen," "what is the forecast," "what if..." "what's next," "what is the plan," "tell this story," and the like. For example, a question "what does this insight mean?" can initiate various follow-up analysis on the datasets or details used to generate the insight, such as descriptions of the formulae, rationales, and data sources used to generate the insight. The formulae can include mathematical or analytic functions used in processing the target
25 datasets to generate final insight objects or intermediate steps thereof. The rationales can include a brief description on why the insight was relevant or chosen for the user, as well as why various formulae, graph types, data ranges, or other properties of the insight object were established. For example, data analysis preferences derived from metadata, initial queries, or past data analysis might indicate that bar chart types are preferred for the
30 datasets.

[0052] Forecasting questions can be queried by the user, such as in the form of "what if" questions related to changing data points, portions of datasets, graph properties, time properties, or other changes. Also, iterative and feedback-generated forecasting can be established where users can select targets for data conclusions or datasets to meet and

examining what data changes would be required to hit the selected targets, such as sales targets or manufacturing targets. These "what if" scenarios can be automatically generated based on the insight datasets, metadata, or queries. Moreover, the insight object can act as a 'model' with which a user can alter parameters, inputs, and properties to see how outputs are affected and predictions are changed.

[0053] Insight results/objects can comprise dynamic insight summaries, verbiage, or data conclusions. These insight summaries can be established as insight objects that explain a key takeaway or key result of another insight object. For example, an insight summary can indicate "sales of model 2.0 were up 26% in Q3 surpassing model 1.0." This summary may be dynamic and tied to the dataset/metadata associated with the insight object, so that when data values or data points change for an insight object, the summary can responsively change accordingly. Data summaries can be provided with the insight results and include titles, graph axis labels, or other textual descriptions of insight objects. The summaries can also include predictive or prospective statements, such as data forecasts over predetermined timeframes, or other statements that are dynamic and change with the insight object.

[0054] For further examples on metadata handling, such as determination and extraction of metadata for various datasets, Figure 3 is presented. Figure 3 includes flow diagram 300 which illustrates an example operation of the elements of Figure 1. In Figure 3, a metadata manager 302 is presented as an example of metadata handler 123. The metadata manager 302 can interface with one or more storage elements (e.g., storage 304), over storage interfaces (e.g., storage interface 314). The storage elements can be examples of cache 132 in Figure 1, although further configurations can be employed. The storage elements can store metadata and user datasets for use during processing by various insight determination elements or recommendation modules, or for usage in later insight requests from users.

[0055] Turning now to the operation of elements of Figure 3, datasets, query information, and user-provided metadata can be delivered to an insight platform that includes a metadata manager 302. The metadata manager 302 can then process the provided datasets/queries/metadata to determine further metadata associated with the datasets. This metadata can be employed in insight processing by one or more recommendation modules. As shown in Figure 3, metadata manager 302 can provide various services such as data type inference, data measure/dimension classification, and data aggregate function detection. Outputs from these services can be provided to a

dataset metadata generation service for processing and output of metadata for the associated datasets.

[0056] A further discussion of the metadata operation continues below. In an example, operation of metadata components illustrated in Figure 3 may comprise the

5 following:

- Metadata is computed once and reused across different recommenders of the insight service.
- Internal subcomponents of the metadata system are typically not recomputed metadata properties that are computed by other
- 10 subcomponents.
- Metadata is cached and typically not recomputed across multiple requests for the metadata.
- Whenever a property of the metadata changes (for example, through a user action), only the metadata properties that depend on the changed property
- 15 are typically recomputed.
- The metadata service can be divided into two major parts. A set of components that compute individual pieces of metadata, and a manager 302 class which holds references to each of the components.

[0057] As mentioned above, various components form the metadata services. The

20 type inference component 306 determines the type of each column of a dataset. A measure v/s dimension classification component 308 classifies each column as a dimension or a measure. An aggregation function detector component 310 suggests aggregation functions for each column. A DatasetMeta generator component 312 generates the DatasetMeta object. A sequential detector component determines whether

25 the data in a column is sequential in nature. It should be noted that the term 'column' can instead be referred to as a 'field' in further examples.

[0058] The metadata manager 302 can comprise a software component 'class' that maintains a list of metadata components. Additionally, the manager 302 class may also maintain an interface to a cache to ensure that re-computation of the metadata for the same

30 input is not necessary. The cache may store a task for every metadata operation being run. This is so that multiple components requesting the metadata can wait on the task if it is still running or directly get the results without waiting if the task has completed. In some examples, the recommenders/providers may only be able to access the metadata through

the manager 302 class.

[0059] An example metadata manager 302 class can be defined as follows:

```

static class MetadataManager
{
5      static MetadataManager()
      {
      }
      public static IMetadata GetMetadata(ITableView data)
      {
10     }
    }
    public interface IMetadata
    {
        Task<IColumnTypes> ColumnTypes { get; }
15     Task<IColumnMeasureDimensionHints> MeasureDimensionHints { get; }
        Task<IColumnAggregationFunctionTypes>
        ColumnAggregationFunctionTypes { get; }
        Task<IColumnSequentialities> ColumnSequentialities { get; }
        Task<DatasetMeta> DatasetMeta { get; }
20     }

```

[0060] Input to each of the metadata processing components can be the raw datasets and any additional metadata that is obtained from the client (for example, cell formats). The metadata components may be aware of the metadata manager 302 so that they can obtain any additional metadata. For example, if the measure/dimension classifier requires

25 column types, it can request types from the manager 302 class which may subsequently call the type detection component, if those types do not already exist in its cache. Each of the components may implement task-based parallelism. This allows multiple components to wait on the results of a component.

[0061] The type inference component 306 may comprise a platform into which

30 multiple type inference providers can be plugged. The provider may accept a standard input and provide types in a standard output format. The input may be a structured form of the data and the output may be a collection of types. Each of the types may have one or more confidence metrics associated with them. The collections of the types from all providers may be provided as input to an aggregation algorithm which may be used to

determine a final type for each column.

[0062] Turning to a further discussion of the elements of Figure 3, the measure/dimension classifier component 308 takes as input the output of the type inference process. The classifier may have a design similar to the type inference system where there may be multiple providers which output their results into an aggregation algorithm to determine the final type decision for one or more columns. The Aggregation Function Detector component 310 generates a list of aggregation functions for measures. The DatasetMeta Generator component 312 creates the DatasetMeta object. The Sequential Data Detector determines whether the given data is sequential in nature.

[0063] Input and Output Interfaces can also be defined for the metadata components. The input to the metadata manager 302 and its components may comprise a form of an interface IRangeData which provides the Cell Values, Cell formats and the Column Headers. The metadata manager 302 and all its components may be agnostic of the column orientation. The metadata manager 302 may detect table orientation in the table recognition step which is independent of metadata detection.

[0064] An example table recognition process can be as follows:

```

interface ITableView
{
    IEnumerable<string> ColumnHeaders { get; }
    IEnumerable<IEnumerable<string>> ColumnData { get; }
    IEnumerable<string> ColumnFormats { get; }
    IEnumerable<IEnumerable<string>> CellFormats { get; }
}
interface IColumnTypes
{
    IEnumerable<FieldType> ColumnTypes { get; }
}
interface IColumnMeasureDimensionHints
{
    IEnumerable<MeasureDimensionHint> MeasureDimensionHints { get; }
}
interface IColumnAggregationFunctionTypes
{
    IEnumerable<IEnumerable<AggrFunc>> AggregationFunctions { get; }
}

```

}

[0065] The internal structure of the type inference component 306 may also be implemented as a platform. Two or more type inference algorithms can be used. A first type inference algorithm may be based on number formatting that is obtained from a client application. A second type inference algorithm may be based on a preprocessor. Each algorithm may take as input a string array representing a single column, and return an array of types for the column. Each type may have a confidence level associated with it. In some examples, the confidence levels may be fed into an aggregation algorithm which may then generate a single type for each column. These types may be added to the DatasetMeta that is passed in. Further examples can add the entire list of types inferred along with the confidence metrics in the DatasetMeta. The internal structure of the dimension/measure classifier component 308 may have a similar pattern as the type inference component 306 with multiple classifiers whose results may be fed into an aggregation algorithm to generate a set of dimensions and a set of measures.

[0066] Further examples of metadata handling components that may be incorporated for generating insights and selecting appropriate insight types for datasets can include implementing a cache so that metadata does not need to be recomputed across multiple requests, and implementing a dependency graph so that on changes to metadata properties, only properties that depend on the changed properties need to be recomputed.

[0067] Figure 4 is a first exemplary method 400 for providing insight results in a productivity application. The method 400 begins at a start operation and flow continues to operation 402 where a dataset and a user query relating to the dataset are received. In some examples the dataset may comprise a plurality of values comprised in one or more columns or rows of a productivity application. In additional examples the dataset may comprise a table or a pivot table of a productivity application. In still other examples, the dataset may comprise a plurality of values obtained from a data source accessed by one or more components of a user platform, such as user platform 110 illustrated in Figure 1, and/or one or more components of an insight platform, such as insight platform 120 illustrated in Figure 1.

[0068] In some examples, the user query received at operation 402 may comprise a natural language question posed by a user of a productivity application. In examples, the user may provide the query to the productivity application via a verbal or typed input type. In other examples, the user query may be initiated by a user providing an input to a productivity application (e.g., hovering a mouse, providing a mouse click, touching a

touch-sensitive display) in the vicinity of a target dataset in the productivity application. Upon receiving the initiation of the user query via the user input to the productivity application, one or more selectable user interface elements may be provided for sending a corresponding user query corresponding to the selected target dataset to one or more components of insight platform 120. In some examples, the selectable user interface elements may be provided for selection based on past user data related to the productivity application and/or past user data related to dataset queries provided to the productivity application.

[0069] From operation 402 flow continues to operation 404 where the dataset is processed to determine metadata that describes one or more properties of the dataset. The metadata may be provided by the user and/or a productivity application associated with the dataset. In examples, the metadata may comprise properties or descriptions associated with the received dataset, such as column and/or row headers, footers, data contexts, data orientations, and application properties of the productivity application. In some examples, the metadata may be determined by a metadata handler to establish metadata for the dataset. For example, a metadata handler may analyze one or more features associated with dataset, such as data features included in the dataset, value types included in the dataset, symbols in the dataset, values included in the dataset, and/or patterns included in the dataset, and assign metadata to the dataset based on the analysis. In some examples, the metadata associated with the dataset may be cached for later processing of the received dataset or datasets that are determined to be similar to the received dataset.

[0070] From operation 404 flow continues to operation 406 where the dataset, metadata, and query are provided to one or more modular recommendation elements for processing into an insight result for the dataset that indicates a result from data analysis directed to the query. The one or more modular recommendation elements may utilize one or more of: past user activity, application usage modalities, organizational traditions with regard to data analysis, and/or individualized data processing techniques, in processing the dataset, metadata, and query. For example, if past user activity associated with the productivity application indicate that the user prefers that one or more specific insight types (e.g., a graph of a dataset, a textual explanation of information associated with a dataset, projections associated with a dataset) be provided based on a query type that is similar to the received query, and/or a dataset type that is similar to the received dataset, the one or more modular recommendation elements may process the dataset, metadata, and query into an insight result corresponding to the user's preferences.

[0071] From operation 406 flow continues to operation 408 where insight results are transferred for use by the productivity application in displaying one or more insight objects based on the insight result. The one or more insight objects may comprise charts, tables, pivot tables, graphs, textual information, interactive visual application elements, selectable application elements for audibly communicating information associated with the dataset, and/or pictures. The one or more insight objects may provide visual and/or audible indications of: information associated with the dataset, summaries of key takeaways associated with the dataset, comparisons of information from the dataset with one or more other datasets related to the dataset, and projections for one or more values or categories associated with dataset.

[0072] In some examples, the one or more values of a dataset corresponding to one or more of the displayed insight objects, and/or metadata associated with a dataset corresponding to one or more of the displayed insight objects, may be interacted with and a display element associated with the interaction may be reflected in one or more affected insight objects. In other examples, one or more of the displayed insight objects may be interacted with, and a corresponding one or more values of the dataset, or a related dataset, may be caused to be modified in associated with the interaction. In additional examples a user may provide, via the productivity application, follow-up queries related to the insight results (e.g., “what happened”, “why did this happen”, “what is the forecast”, “what if . . .”, “what’s next”, “what is the plan”, “tell this story”), and additional analysis may be performed for providing information related to a received follow-up query (e.g., providing a description of formulae utilized in generating the insight results, providing a description of rationales for the displayed insight objects, providing a description of data sources used to generated the displayed insight objects).

[0073] From operation 408 the method 400 continues to an end operation, and the method 400 ends.

[0074] Figure 5 is a second exemplary method 500 for providing dataset insights for a productivity application. The method 500 begins at a start operation and flow continues to operation 502 where an indication to generate an insight associated with a dataset is received. The indication may comprise a typed command, a verbal command, a command issued via a mouse click, a command issued by interacting with the dataset, a user interaction associated with a user interface element of a productivity application, and/or an automatic indication received based on automated analysis of one or more datasets associated with a productivity application (e.g., an analysis of one or more datasets based

on the datasets being created, the analysis of one or more datasets based on information associated with the one or more datasets being modified).

[0075] From operation 502 flow continues to operation 504 where one or more properties associated with the dataset are analyzed. The one or more properties may
5 comprise values include in the dataset, values of one or more datasets related to the dataset, column headers associated with the dataset, column footers associated with the dataset, font properties of data in the dataset, relationships of data in the dataset to one or more other datasets, and metadata associated with the dataset. According to some
10 examples, the analysis of the one or more properties may comprise identifying one or more patterns associated with a plurality of values in the dataset, identifying relationships of the dataset to one or more other datasets, and identifying past user interaction related to the dataset or one or more similar datasets.

[0076] From operation 504 flow continues to operation 506 where a category type is assigned to a plurality of values of the dataset based on the analysis of the one or more
15 properties at operation 504. In examples, the category type may comprise a value type (e.g., a text value type, a number value type, a symbol value type, a denomination value type, a date value type, a specific function value type, an address value type, a person name value type, and an object type value type (e.g., company names, book names, social security numbers, performance ratings, sales figures, geographic locations, colors, shapes,
20 category types).

[0077] From operation 506 flow continues to operation 508 an insight associated with the dataset is generated by applying at least one function to a plurality of values of the dataset. In examples, the at least one function may comprise one or more of: a sort
function, an averaging function, an add function, a subtract function, a multiply function, a
25 divide function, a graph generation function, a chart generation function, a pattern identification function, a summarization function, and a projection function. In examples, the at least one function may be applied based on past user history associated with the productivity application, a type of user query corresponding to the received indication to generate the insight, and the ability to apply the at least one function to value types
30 included in the dataset.

[0078] From operation 508 flow continues to operation 510 where the generated insight is caused to be displayed in a user interface of the productivity application. In examples, the displayed insight may comprise charts, tables, pivot tables, graphs, textual information, interactive visual application elements, selectable application elements for

audibly communicating information associated with the dataset, and/or pictures. The displayed insight may provide visual and/or audible indications of: information associated with the dataset, summaries of key takeaways associated with the dataset, comparisons of information from the dataset, summaries of key takeaways associated with the dataset, comparisons of information of information from the dataset with one or more other datasets related to the dataset, and projections for one or more values or categories associated with the dataset.

[0079] From operation 510 flow continues to an end operation, and the method 500 ends.

[0080] The systems, methods, and devices described herein provide technical advantages for interacting and viewing information associated with productivity applications. For example, users may be provided with dataset insights, which may be generated with a specific querying user taken into account, that visually and/or audibly communicate key takeaways associated with a dataset, summaries of information included in a dataset, comparisons of data in a dataset, comparisons of data in a dataset with data from other related datasets, and projections associated with a dataset.

[0081] As described herein, an insight service may process dataset insight queries in a single, portable, format via an insight API, and provide one or more generated insights of one or more insight types, to a plurality of different application types (which may each support various different insight features), in a portable format. The ability of the insight service to uniformly analyze, process, and generate insights in a portable format reduces processing costs (i.e., CPU cycles) that would otherwise be required for multiple application-specific insight services, or multiple application-specific insight service engines, to perform that analysis, processing, and generation of insights that are specific to each application type from which insight queries may be received.

[0082] The ability to generate insights for datasets based on the analysis of user provided metadata for datasets, metadata associated with datasets based on dataset creation, and/or the association of metadata with datasets based on analysis of dataset information via an insight service and the mechanisms described herein allows for the surfacing of summary and/or key information associated with datasets, which can be interacted with in various ways in order to quickly view the result of modifications to surfaced insights and/or dataset values. These enhanced features provide a better user experience, the ability to quickly and efficiently identify and view relevant information associated with large datasets, which may not otherwise be readily identifiable due to the

size of a dataset, and cost savings at least in the time required to identify relevant data in productivity applications, and the processing costs required to identify relevant data in datasets and navigate large datasets comprised in productivity applications and/or datasets from which one or more values of a productivity application depend.

5 **[0083]** Turning now to Figure 6, computing system 601 is presented. Computing system 601 that is representative of any system or collection of systems in which the various operational architectures, scenarios, and processes disclosed herein may be implemented. For example, computing system 601 can be used to implement any of user platform 110 or insight platform 120 of Figure 1. Examples of computing system 601
10 include, but are not limited to, server computers, cloud computing systems, distributed computing systems, software-defined networking systems, computers, desktop computers, hybrid computers, rack servers, web servers, cloud computing platforms, and data center equipment, as well as any other type of physical or virtual server machine, and other computing systems and devices, as well as any variation or combination thereof. When
15 portions of computing system 601 are implemented on user devices, example devices include smartphones, laptop computers, tablet computers, desktop computers, gaming systems, entertainment systems, and the like.

[0084] Computing system 601 may be implemented as a single apparatus, system, or device or may be implemented in a distributed manner as multiple apparatuses, systems, or
20 devices. Computing system 601 includes, but is not limited to, processing system 602, storage system 603, software 605, communication interface system 607, and user interface system 608. Processing system 602 is operatively coupled with storage system 603, communication interface system 607, and user interface system 608.

[0085] Processing system 602 loads and executes software 605 from storage system
25 603. Software 605 includes insights environment 606, which is representative of the processes discussed with respect to the preceding Figures. When executed by processing system 602 to enhance data insight generation and handling, software 605 directs processing system 602 to operate as described herein for at least the various processes, operational scenarios, and environments discussed in the foregoing implementations.
30 Computing system 601 may optionally include additional devices, features, or functionality not discussed for purposes of brevity.

[0086] Referring still to Figure 6, processing system 602 may comprise a microprocessor and processing circuitry that retrieves and executes software 605 from storage system 603. Processing system 602 may be implemented within a single

processing device, but may also be distributed across multiple processing devices or sub-systems that cooperate in executing program instructions. Examples of processing system 602 include general purpose central processing units, application specific processors, and logic devices, as well as any other type of processing device, combinations, or variations thereof.

[0087] Storage system 603 may comprise any computer readable storage media readable by processing system 602 and capable of storing software 605. Storage system 603 may include volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information, such as computer readable instructions, data structures, program modules, or other data. Examples of storage media include random access memory, read only memory, magnetic disks, resistive memory, optical disks, flash memory, virtual memory and non-virtual memory, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other suitable storage media. In no case is the computer readable storage media a propagated signal.

[0088] In addition to computer readable storage media, in some implementations storage system 603 may also include computer readable communication media over which at least some of software 605 may be communicated internally or externally. Storage system 603 may be implemented as a single storage device, but may also be implemented across multiple storage devices or sub-systems co-located or distributed relative to each other. Storage system 603 may comprise additional elements, such as a controller, capable of communicating with processing system 602 or possibly other systems.

[0089] Software 605 may be implemented in program instructions and among other functions may, when executed by processing system 602, direct processing system 602 to operate as described with respect to the various operational scenarios, sequences, and processes illustrated herein. For example, software 605 may include program instructions for implementing the dataset processing environments and platforms discussed herein.

[0090] In particular, the program instructions may include various components or modules that cooperate or otherwise interact to carry out the various processes and operational scenarios described herein. The various components or modules may be embodied in compiled or interpreted instructions, or in some other variation or combination of instructions. The various components or modules may be executed in a synchronous or asynchronous manner, serially or in parallel, in a single threaded environment or multi-threaded, or in accordance with any other suitable execution

paradigm, variation, or combination thereof. Software 605 may include additional processes, programs, or components, such as operating system software or other application software, in addition to or that include insights environment 606. Software 605 may also comprise firmware or some other form of machine-readable processing instructions executable by processing system 602.

[0091] In general, software 605 may, when loaded into processing system 602 and executed, transform a suitable apparatus, system, or device (of which computing system 601 is representative) overall from a general-purpose computing system into a special-purpose computing system customized to facilitate data insight generation and handling.

Indeed, encoding software 605 on storage system 603 may transform the physical structure of storage system 603. The specific transformation of the physical structure may depend on various factors in different implementations of this description. Examples of such factors may include, but are not limited to, the technology used to implement the storage media of storage system 603 and whether the computer-storage media are characterized as primary or secondary storage, as well as other factors.

[0092] For example, if the computer readable storage media are implemented as semiconductor-based memory, software 605 may transform the physical state of the semiconductor memory when the program instructions are encoded therein, such as by transforming the state of transistors, capacitors, or other discrete circuit elements constituting the semiconductor memory. A similar transformation may occur with respect to magnetic or optical media. Other transformations of physical media are possible without departing from the scope of the present description, with the foregoing examples provided only to facilitate the present discussion.

[0093] Insights environment 606 includes one or more software elements, such as OS 621 and applications 622. These elements can describe various portions of computing system 601 with which users, dataset sources, machine learning environments, or other elements, interact. For example, OS 621 can provide a software platform on which application 622 is executed and allows for processing datasets for insights and visualizations, among other functions. In one example, insight processor 623 implements elements from insight platform 120 of Figure 1, namely elements 122-124.

[0094] Communication interface system 607 may include communication connections and devices that allow for communication with other computing systems (not shown) over communication networks (not shown). Examples of connections and devices that together allow for inter-system communication may include network interface cards, antennas,

power amplifiers, RF circuitry, transceivers, and other communication circuitry. The connections and devices may communicate over communication media to exchange communications with other computing systems or networks of systems, such as metal, glass, air, or any other suitable communication media. Physical or logical elements of communication interface system 607 can receive datasets, transfer datasets, metadata, and control information between one or more distributed data storage elements, and interface with a user to receive data selections and provide insight results, among other features.

[0095] User interface system 608 is optional and may include a keyboard, a mouse, a voice input device, a touch input device for receiving input from a user. Output devices such as a display, speakers, web interfaces, terminal interfaces, and other types of output devices may also be included in user interface system 608. User interface system 608 can provide output and receive input over a network interface, such as communication interface system 607. In network examples, user interface system 608 might packetize display or graphics data for remote display by a display system or computing system coupled over one or more network interfaces. Physical or logical elements of user interface system 608 can receive datasets or insight selection information from users or other operators, and provide processed datasets, insight results, or other information to users or other operators. User interface system 608 may also include associated user interface software executable by processing system 602 in support of the various user input and output devices discussed above. Separately or in conjunction with each other and other hardware and software elements, the user interface software and user interface devices may support a graphical user interface, a natural user interface, or any other type of user interface.

[0096] Communication between computing system 601 and other computing systems (not shown), may occur over a communication network or networks and in accordance with various communication protocols, combinations of protocols, or variations thereof. Examples include intranets, internets, the Internet, local area networks, wide area networks, wireless networks, wired networks, virtual networks, software defined networks, data center buses, computing backplanes, or any other type of network, combination of network, or variation thereof. The aforementioned communication networks and protocols are well known and need not be discussed at length here.

However, some communication protocols that may be used include, but are not limited to, the hypertext transfer protocol (HTTP), Internet protocol (IP, IPv4, IPv6, etc.), the transmission control protocol (TCP), and the user datagram protocol (UDP), as well as any

other suitable communication protocol, variation, or combination thereof.

[0097] The functional block diagrams, operational scenarios and sequences, and flow diagrams provided in the Figures are representative of exemplary systems, environments, and methodologies for performing novel aspects of the disclosure. While, for purposes of simplicity of explanation, methods included herein may be in the form of a functional diagram, operational scenario or sequence, or flow diagram, and may be described as a series of acts, it is to be understood and appreciated that the methods are not limited by the order of acts, as some acts may, in accordance therewith, occur in a different order and/or concurrently with other acts from that shown and described herein. For example, those skilled in the art will understand and appreciate that a method could alternatively be represented as a series of interrelated states or events, such as in a state diagram. Moreover, not all acts illustrated in a methodology may be required for a novel implementation.

[0098] The descriptions and figures included herein depict specific implementations to teach those skilled in the art how to make and use the best option. For the purpose of teaching inventive principles, some conventional aspects have been simplified or omitted. Those skilled in the art will appreciate variations from these implementations that fall within the scope of the disclosure. Those skilled in the art will also appreciate that the features described above can be combined in various ways to form multiple implementations. As a result, the invention is not limited to the specific implementations described above, but only by the claims and their equivalents.

CLAIMS

1. A method of providing insight results in a productivity application, the method comprising:
 - receiving a dataset and a user query relating to the dataset;
 - processing the dataset to determine metadata that describes one or more properties of the dataset;
 - providing the dataset, the determined metadata, and the user query to one or more modular recommendation elements for processing into an insight result for the dataset, wherein the insight result indicates a result from data analysis directed to the query; and
 - transferring the insight result for use by the productivity application in displaying one or more insight objects based on the insight result.
2. The method of claim 1, wherein a specific function is selected, from a plurality of functions, for processing a plurality of values of the dataset based on the specific function having a same property type assigned to it as a property type described in metadata for the plurality of values of the dataset.
3. A system for providing dataset insights for a productivity application, the system comprising:
 - a memory for storing executable program code; and
 - one or more processors, functionally coupled to the memory, the one or more processors being responsive to computer-executable instructions contained in the program code and operative to:
 - receive an indication to generate an insight associated with a dataset;
 - analyze one or more properties associated with the dataset;
 - assign, based on analyzing the one or more properties, a category type to a plurality of values of the dataset;
 - generate an insight associated with the dataset, wherein the insight is generated by applying at least one function to a plurality of values of the dataset; and
 - cause the generated insight to be displayed in a user interface of the productivity application.
4. The system of claim 3, wherein the applied function is selected from a plurality of functions, each of the plurality of functions being associated with one or more value category types, and wherein a function of the plurality of functions can be applied to a value that has a same category type assigned to it as a corresponding function.
5. The system of claim 3, wherein analyzing the one or more properties comprises

identifying one or more patterns associated with a plurality of values in the dataset.

6. The system of claim 3, wherein an insight type for the displayed insight is selected from one of: a graph associated with a plurality of values of the dataset; a chart associated with a plurality of values of the dataset; and a pivot table associated with a plurality of values of the dataset.

7. The system of claim 3, wherein the one or more processors are further responsive to the computer-executable instructions contained in the program code and operative to:

assign a category type to a plurality of values in the dataset based on the analysis of the one or more properties of the dataset;

identify a second dataset associated with the assigned category type; and

utilize one or more values included in the second dataset in at least one of the applied functions.

8. The system of claim 7, wherein the one or more processors are further responsive to the computer-executable instructions contained in the program code and operative to:

receive an indication to provide information associated with one or more criteria utilized in generating the insight; and

cause, based on the received indication, a description of the second dataset to be displayed in the user interface of the productivity application, wherein the description of the second dataset comprises one of: a location in the productivity application where the second dataset is located; a file where the second dataset is located; a remote storage location where the second dataset is located; and a web address where the second dataset is located.

9. The system of claim 8, wherein the one or more processors are further responsive to the computer-executable instructions contained in the program code and operative to:

cause metadata associated with the dataset, and historical user data for the spreadsheet application to be cached;

receive an indication to generate a second insight associated with the dataset; and

cause the second insight to be displayed in a user interface of the productivity application, in an insight type specific to the cached metadata and the cached historical user data.

10. The system of claim 3, wherein the one or more processors are further responsive to the computer-executable instructions contained in the program code and operative to:

receive a user input to modify a display property of the displayed insight;

identify a value of the dataset utilized in at least one of the applied functions that

resulted in the display of the display property in the displayed insight;

determine a new value for the identified value that will result in modification of the display property when utilized in the at least one applied function; and

cause the new value to be displayed in the user interface of the productivity application.

11. The system of claim 10, wherein the display property comprises one of: a size of a displayed metric for a category in a chart of the displayed insight; a color of a displayed metric for a category in a chart of the displayed insight; a size of a displayed metric for a category in a graph of the displayed insight; and a color of a displayed metric for a category in a graph of the displayed insight.

12. The system of claim 3, wherein the one or more processors are further responsive to the computer-executable instructions contained in the program code and operative to:

receive a modification to a value in the dataset utilized in generating the insight;
re-apply at least one of the applied functions to the value based on the received modification; and

modify one or more display properties of the displayed insight in the user interface of the productivity application based on the result of the re-applied at least one function.

13. A computer-readable storage device comprising executable instructions that, when executed by one or more processors, assists with providing dataset insights for a productivity application, the computer-readable storage device including instructions executable by the one or more processors for:

receiving a user query to generate an insight associated with a dataset;
causing, based on an assigned value category type associated with one or more values of the dataset, a plurality of insight display type categories selectable for generating insight results for the dataset to be displayed in a user interface of the productivity application;

receiving a selection of one of the insight display type categories; and
causing an insight result for the dataset corresponding to the selected insight display type category to be displayed in the user interface of the productivity application.

14. The computer-readable storage device of claim 13, wherein the value category type associated with the one or more values of the dataset is assigned based on at least one of: data included in one or more column headers associated with the dataset; data included in one or more column footers associated with the dataset; and user-provided data that describes the one or more values of the dataset.

15. The computer-readable storage device of claim 13, wherein the instructions are further executable by the one or more processors for:

receiving an API call-in from the productivity application for the insight; and
sending, based on the API call-in, the user query, and the assigned value category type associated with the one or more values of the dataset, the insight result to the productivity application in a standardized output for interpretation by a plurality of application types.

100





FIG. 1

2/6

FIG. 2






FIG. 3

4/6

FIG. 4

400

5/6

FIG. 5

500




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/062154

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F16/26
ADD.

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)
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, INSPEC, 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	Anonymous: "Watson Analytics Review: Natural Language Query IBM Blueview", 26 April 2015 (2015-04-26), XP055551393, Retrieved from the Internet: URL:http://ibmbblueview.com/watson-analytic s-review-natural-language-query/ [retrieved on 2019-02-04] section "Ask And Ye Shall Receive"; pages 1-5	1-15
X	US 2013/145244 A1 (ROTHSCHILLER CHAD BARRY [US] ET AL) 6 June 2013 (2013-06-06) paragraph [0004] ----- -/-	1-15



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

6 February 2019

Date of mailing of the international search report

14/02/2019

Name and mailing address of the ISA/

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

Authorized officer

Correia Martins, F

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/062154

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2015/370872 A1 (RAGHAVAN SRINIVASAN SUNDAR [US] ET AL) 24 December 2015 (2015-12-24) paragraph [0024] -----	1-15
X	DASORIYA RAYAN: "A review of big data analytics over cloud", 2017 IEEE INTERNATIONAL CONFERENCE ON CONSUMER ELECTRONICS-ASIA (ICCE-ASIA), IEEE, 5 October 2017 (2017-10-05), pages 1-6, XP033327950, DOI: 10.1109/ICCE-ASIA.2017.8307833 [retrieved on 2018-03-07] abstract -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2018/062154

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2013145244	A1	06-06-2013	
		AU 2012348174 A1	26-06-2014
		BR 112014013454 A2	13-06-2017
		CA 2857409 A1	13-06-2013
		CN 102982016 A	20-03-2013
		EP 2788852 A1	15-10-2014
		JP 2015507779 A	12-03-2015
		KR 20140105733 A	02-09-2014
		RU 2014122851 A	10-12-2015
		US 2013145244 A1	06-06-2013
		WO 2013085727 A1	13-06-2013

US 2015370872	A1	24-12-2015	
		US 2015370872 A1	24-12-2015
		US 2015370881 A1	24-12-2015
		US 2015370882 A1	24-12-2015
		US 2015370883 A1	24-12-2015
