

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(10) International Publication Number

WO 2016/049034 A1

(43) International Publication Date
31 March 2016 (31.03.2016)

(51) International Patent Classification:
G06F 17/30 (2006.01)

(21) International Application Number:
PCT/US2015/051462

(22) International Filing Date:
22 September 2015 (22.09.2015)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/054,517 24 September 2014 (24.09.2014) US
14/678,218 3 April 2015 (03.04.2015) US

(71) Applicant: ORACLE INTERNATIONAL CORPORATION [US/US]; 500 Oracle Parkway; Mail Stop 50P7, Redwood Shores, CA 94067 (US).

(72) Inventors: SHEFFER, Uri; 155 Coolidge Street, Apt. 1, Brookline, MA 02446 (US). POCOCK, Adam, Craig; 232 Arboretum Way, Burlington, MA 01803 (US). STEVENS, Brook; 89 Varnum Street, Arlington, MA 02474 (US). ISHAQUE, Mashhood; 57 Randolph Street, Arlington, MA 02474 (US). ZELEVINSKY, Vladimir; 18 Fales Road, Sharon, MA 02142 (US). SPAULDING, Tristan, R.; 22 Quincy Street, Medford, MA 02155 (US).

(74) Agents: GOLDSMITH, Barry, S. et al.; Miles & Stockbridge, P.C., 1751 Pinnacle Drive, Suite 1500, Tysons Corner, VA 22102-3833 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, 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))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) Title: GUIDED DATA EXPLORATION




Fig. 1

(57) Abstract: A system for exploring data receives the data from a database and indexes the data in a server. The system displays one or more selectable datasets from the indexed data, where the selected dataset includes one or more attributes. The system then sorts the one or more attributes by level of interestingness and displays the sorted attributes.

GUIDED DATA EXPLORATION

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority of Provisional Patent Application Serial No. 62/054,517, filed on September 24, 2014, the contents of which is hereby incorporated by reference.

FIELD

[0002] One embodiment is directed generally to a computer system, and in particular to a computer system that provides analysis of data through guided data exploration.

BACKGROUND INFORMATION

[0003] In recent years, computer systems have enabled individuals and organizations to capture and store vast quantities of data. The proliferation of data is sometime referred to as “big data”, which is an all-encompassing term for any collection of datasets so large or complex that it becomes difficult to process using traditional data processing applications.

[0004] The challenges of big data include analysis, capture, curation, search, sharing, storage, transfer, visualization, and privacy concerns. The existence of such large quantities of data has led to an ever increasing need for improved systems and methods of analyzing and exploring data.

SUMMARY

[0005] One embodiment is a system for exploring data. The system receives the data from a database and indexes the data in a server. The system displays one or more selectable datasets from the indexed data, where the selected dataset includes one or more attributes. The system then sorts the one or more attributes by level of interestingness and displays the sorted attributes.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] Fig. 1 is a block diagram of a computer server/system in accordance with an embodiment of the present invention.

[0007] Fig. 2 is a flow diagram of the functionality of the guided data exploration module of Fig. 1 and other elements in accordance with one embodiment of the present invention.

[0008] Fig. 3 illustrates an example user interface displaying available datasets in accordance with one embodiment.

[0009] Fig. 4 illustrates an example user interface displaying the attributes shown as sorted “tiles” after the selection of the “wine sales” dataset of Fig. 3 in accordance with an embodiment.

[0010] Fig. 5 illustrates a graph of the interestingness in relation to the normalized entropy in accordance with one embodiment.

[0011] Fig. 6 is a flow diagram of the functionality of the guided data exploration module of Fig. 1 and other elements in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION

[0012] Embodiments of the invention provide guided data exploration. One of the embodiments generates a user interface that displays indexed datasets (i.e., collections of data or data records) and allows a user to select one of the datasets. In response, the user interface displays attributes of the selected dataset in order of interestingness. The user can then further select one or more values of the attributes, which results in a refining of the dataset to assist the user in the selection and display of the desired data.

[0013] Fig. 1 is a block diagram of a computer server/system 10 in accordance with an embodiment of the present invention. Although shown as a single system, the functionality of system 10 can be implemented as a distributed system. Further, the functionality disclosed herein can be implemented on separate servers or devices that may be coupled together over a network. Further, one or more component of system 10 may not be included. For example, for functionality of a user client, system 10 may be a smartphone that includes a processor, memory and a display, but may not include one or more of the other components shown in Fig. 1.

[0014] System 10 includes a bus 12 or other communication mechanism for communicating information, and a processor 22 coupled to bus 12 for processing information. Processor 22 may be any type of general or specific purpose processor. System 10 further includes a memory 14 for storing information and instructions to be executed by processor 22. Memory 14 can be comprised of any combination of random access memory (“RAM”), read only memory (“ROM”), static

storage such as a magnetic or optical disk, or any other type of computer readable media. System 10 further includes a communication device 20, such as a network interface card, to provide access to a network. Therefore, a user may interface with system 10 directly, or remotely through a network, or any other method.

[0015] Computer readable media may be any available media that can be accessed by processor 22 and includes both volatile and nonvolatile media, removable and non-removable media, and communication media. Communication media may include computer readable instructions, data structures, program modules, or other data in a modulated data signal such as a carrier wave or other transport mechanism, and includes any information delivery media.

[0016] Processor 22 is further coupled via bus 12 to a display 24, such as a Liquid Crystal Display (“LCD”). A keyboard 26 and a cursor control device 28, such as a computer mouse, are further coupled to bus 12 to enable a user to interface with system 10.

[0017] In one embodiment, memory 14 stores software modules that provide functionality when executed by processor 22. The modules include an operating system 15 that provides operating system functionality for system 10. The modules further include a guided data exploration module 16 for providing guided data exploration functionality, and all other functionality disclosed herein. System 10 can be part of a larger system. Therefore, system 10 can include one or more additional functional modules 18 to include the additional functionality, such as an enterprise resource planning (“ERP”) module that generates ERP data to be explored by module 16. A database 17 is coupled to bus 12 to provide centralized storage for modules 16 and 18 and store unstructured data, structured data, etc.

[0018] Each module can include multiple modules. In one embodiment, module 16 includes an indexing module that receives data from a database and indexes the data, a display module that displays one or more selectable datasets from the indexed data, and a sorting module that receives a selection of one or more of the selectable datasets and sorts the one or more attributes by level of interestingness.

[0019] Fig. 2 is a flow diagram of the functionality of guided data exploration module 16 and other elements in accordance with one embodiment of the present invention. In one embodiment, the functionality of the flow diagrams of Fig. 2 and Fig. 6 is implemented by software stored in memory or other computer readable or tangible medium, and executed by a processor. In other embodiments, the functionality may be performed by hardware (e.g., through the use of an application-specific integrated circuit ("ASIC"), a programmable gate array ("PGA"), a field programmable gate array ("FPGA"), etc.), or any combination of hardware and software.

[0020] At 202, unstructured or partially structured data is stored in database 17 of Fig. 1. In one embodiment, the data is stored in an Apache Hive, which is a data warehouse infrastructure built on top of Hadoop for providing data summarization, query, and analysis. However, any type of unstructured or structured data can be received and stored at 202. For example, in one embodiment the data is a large amount of Twitter "tweets". In one embodiment, the data is stored in a Hadoop Distributed File System ("HDFS").

[0021] At 204, metadata discovery is performed on the data from 202, and the data is then enriched according to the results of metadata discovery. In one

embodiment, metadata discovery is performed to determine the nature/type of data from the data itself (e.g., is the data a number, a geo-code, etc.), including the cardinality of the data. In one embodiment, the data discovery is performed by the “Endeca Information Discovery Studio Provisioning Service”, from Oracle Corp. The metadata discovery generates a list of metadata that feeds into a user interface. In one embodiment, the enrichment is not performed. In another embodiment, both the type discovery and enrichment is not performed.

[0022] Table 1 below provides some example metadata:

Name	Description
id	Value that uniquely identifies a dataset.
collectionKey	Name of the collection that contains the data for this dataset.
sourceType	Where this dataset gets its data from.
sourceName	The name of the database/table/file that this dataset gets its data from.
displayName	This will be displayed as the dataset name
description	Description of the dataset.
timesViewed	The number of times this dataset has been viewed by a user.
timesFavorited	The number of times this dataset has been favorited by a user.

Table 1

[0023] At 206, the data from 204 is indexed into a server and published to a user interface. In one embodiment, the data is indexed at 206 as an Endeca index in an “MDEX” engine from Oracle Corp.

[0024] At 208, the user interface is generated that displays all available datasets/data records or catalog of datasets that are indexed on the server. Fig. 3 illustrates an example user interface 300 displaying the available datasets at 301 in accordance with one embodiment. A total of 675 datasets are available via the interface (not all shown in Fig. 3). Each dataset is selectable by the user. In one

embodiment, only a single dataset may be available and shown at 208. A selection of one of the datasets can also be received at 208. In one embodiment, a selection is not required at 208.

[0025] In response to a selection of a dataset at 208, the attributes are sorted by level of interestingness at 210 and an “explorer” user interface is generated at 212 for the selected dataset. The attributes of the dataset are listed for further exploration/selection. Fig. 4 illustrates an example user interface 400 displaying the attributes shown as sorted “tiles” after the selection of the “wine sales” dataset 305 of Fig. 3 in accordance with an embodiment. Example tiles include a “delivery date” tile 401, a “zip code” tile 402, a “date of delivery” tile 403, etc. Each tile summarizes a variable/attribute from the selected data set, and the tiles are sorted by order of interestingness (i.e., a computation of a predictor of how likely the user will find each attribute, and the corresponding representative tile for that attribute, interesting) so that the attributes that most explain the dataset are initially displayed. In one embodiment, as described in detail below, entropy is used to determine the sorting by levels of interestingness. In another embodiment, the representative tile are sorted alphabetically. In another embodiment, the representative tiles are listed in the order that the tiles occur in the data (i.e., no sorting).

[0026] User interface 400 further includes a list of available refinements shown on the left side at 410. As part of the generation of tiles at 212, the type of visualization (e.g., bar chart, graph, map, etc.) is also determined. In one embodiment, a hard coded decision tree is used to determine the type of visualization. In Fig. 4, each tile represents an attribute of a dataset (i.e., a column of database attributes). A user can also zoom into a particular tile/column.

Therefore, a user can have at a glance a view of each particular column/tile that includes the metadata.

[0027] Referring again to Fig. 2, at 214 the user selects (or unselects) from the list of available refinements 410, one or more values from one or more of the attributes, refining the datasets.

[0028] At 216, the attributes/tiles are again sorted by level of interestingness based on the current dataset, and the user interface showing a revised set of selectable data records is generated at 212. 214 and 216 can be repeated as many times as necessary so that the selection of the records from the dataset is iteratively refinable.

[0029] The embodiment of Fig. 2 is considered "univariate" because each column is considered separately.

[0030] In one embodiment, entropy is used to determine the level of interestingness of tiles/attributes at 210 of Fig. 2. "Entropy" is a measurement of the uncertainty in a random variable. The typical unit of measure used with entropy is a "bit". The more uncertain the outcome of the random variable, the more bits are needed to represent the different values.

[0031] As an example, assume there is a need to record the outcome of an experiment that can either be "1000" or "2000". First, since there are only two possible outcomes, there is no need for the entire memory (bit-wise) representation "1000" or "2000". Instead, a convention of "0" to indicate the former and "1" to indicate the latter can be used. The true representation of the outcome is therefore only 1 bit, and forms the upper limit to the entropy of this random variable.

[0032] Second, if the probability of the experiment's outcome is known, the

entropy value can be further diminished, since the uncertainty inherent of this variable has been reduced. For example, tossing an unbiased coin yields an equal 0.5 chance (probability) of a tails or heads outcome. Since the uncertainty is high, the entropy would reflect its highest value (i.e., 1). If, however, the outcome records whether women are pregnant or not, and it is known that pregnant women account for 5% of the women population, the entropy will drop and indicate a value of 0.2864 bits.

[0033] Every attribute of the datasets, such as datasets 301 of Fig. 3, has its own entropy value (also referred to as "self-information") and denotes how much information is given by this attribute. This information (also known as "information gain") can be determined based on the uncertainty of the attribute. For example, if all the values of a particular attribute are the same (a single value outcome), the attribute carries no signal, and therefore has zero information gain.

[0034] The entropy value calculated for a single attribute is a non-negative number and ranges from 0 up to $\log |x|$ (pronounced as "log of count of x"), where $|x|$ is the number of different values of this attribute. For the values of entropy to be expressed in bits, the logarithm is taken in base 2. In such a case, for example, the column with four equally distributed values carries $\log 4 = 2$ bits of information. As a different example, for a variable whose outcome is always a zero, the variable has only one outcome value and carries no information, its entropy value being $\log 1 = 0$. If, to provide another example, x is denoting a uniformly distributed value between 0 and 65536, there is high uncertainty of each value. In this example, there are $65536 = 2^{16}$ outcomes and, hence, 16 bits of entropy.

[0035] Some embodiments compute entropy as Shannon entropy according to

the following formula:

$$H(X) = - \sum_{i=1}^n p(x_i) \log_b p(x_i)$$

where $H(X)$ is the entropy of variable X , index i loops over all possible n outcomes, x_i represents the possible outcome, and $p(x_i)$ is the probability for outcome x_i . In one embodiment, binary (base 2) logarithm is used, in which case the resulting entropy is measured in bits.

[0036] Since all probabilities are given as a value in the range between 0 and 1, all log outcomes are negative, hence the negation outside the summation.

In addition, the lower the probability of an outcome, the smaller the probability value and hence the higher the log value. In other words, the infrequent occurring values are, in fact, the biggest contributors to the entropy value.

[0037] Some embodiments may normalize the entropy by dividing it by $\log |n|$, making the resulting normalized entropy fall in the range 0 – 1.

[0038] In general, entropy can be evaluated for any discrete variable. In the case of numerical variables, the entropy can be either calculated directly (via the computationally-complex differential approach), or the data can be discretized, or binned, to convert it into a categorical variable. The latter approach potentially causes marginal loss of precision, but gains a considerable simplicity and speed of computation.

[0039] In the discrete variable case, the probabilities are the frequencies of the attributes within the dataset. In one embodiment, high cardinality data (e.g., variables that have the number of different values comparable to the number of records in the system; product reviews is an example of such high cardinality data,

since every review can reasonably be expected to be different) can be assumed to be uninteresting. In the case of a variable that contains natural language text, term extraction can be used to convert such high cardinality variable into a lower cardinality (and, thus, more interesting) variable.

[0040] In one embodiment, the use of entropy values includes calculating the entropy of every attribute of a dataset, normalized to a 0 – 1 range. Further uses include sorting of the attributes based on descending order of an outcome of an interestingness function, with an attempt to give prominence to the attributes that are more interesting, as described below.

[0041] In general, what may be interesting for one data analyst (or scientist) may not be so interesting for another. However, there are two degenerate cases that are simply not interesting by their nature. On the low end of the entropy range (closer to 0) are attributes that hold only a single outcome. Intuitively, these type do not contribute to the overall understanding of the data set. Similarly, on the high end of the entropy values (closer to the normalized value of 1), exist high cardinality attributes (e.g., a product inventory number column, which is expected to be different for every product). Such attributes are not expected to carry information either.

[0042] At the values above the low end degenerate case, interesting signal appears, signifying variation on the outcome of this variable. Similarly, there is interesting signal at the values below the high end degenerate case.

[0043] Certain embodiments posit a way to translate the precisely-computed entropy into interestingness, by first applying the low-end and high-end cutoffs and then mapping the extreme (low and high) values of entropy to high interestingness, while mapping intermediate values to low interestingness.

[0044] Fig. 5 illustrates one possible graph of the interestingness in relation to the normalized entropy. In Fig. 5, the x-axis is the normalized entropy from 0 to 1 with 0.5 being the minimum of the mapping curve. The y-axis represents how interesting an attribute will be. In some embodiments the cutoffs for low and high levels of entropy are different and could be tuned independently.

[0045] Since the values of entropy span from 0 to $\log |n|$, where n is the number of different outcomes (or values) for each particular column or attribute, comparing entropy values for the attributes with a different number of outcomes can present a challenge. Some embodiments perform entropy normalization by dividing entropy for each particular column by $\log |n|$. The values of normalized entropy fall between 0 and 1 and thus can be compared directly.

[0046] For example, consider two columns with the same value of entropy 0.5. If the first column contains only values "true" and "false", it has $n = 2$, and the normalized entropy is $0.5 / \log(2) = 0.5$. The same process applied to another column with the same value of entropy 0.5 but with four different values would result in the normalized entropy $0.5 / \log(4) = 0.25$. In this example, two columns with the same value of entropy before normalization would have different values of normalized entropy and, thus, different interestingness.

[0047] Other embodiments may utilize different mappings of entropy to interestingness. For example, a parabolic curve with the global minimum in the 0 – 1 range would also satisfy the above considerations.

[0048] Certain embodiments can apply different mappings of entropy to interestingness based on each attribute type. For example, geocodes can be considered always interesting, no matter the distribution of their values.

[0049] Some embodiments allow users to dynamically modify the lists of the attributes that have been sorted according to their interestingness. The possibilities include user interface elements such as “remove” and “like” buttons, to correspondingly exclude and promote selected attributes.

[0050] Some embodiments add the utilization of machine learning approaches to further determine the specific ranges/thresholds for the degenerate cases based on the interest shown by users.

[0051] The embodiments described above compute entropy for each attribute in isolation. In other embodiments, entropy calculations are performed at the bivariate analysis level instead of the univariate level. This allows entropy to be computed between two different attributes (mutual and conditional entropy). In this embodiment, the user may select an attribute before the interestingness-based sorting. In this scenario, the entropy is computed and the interestingness is determined relative to the selected column.

[0052] In another embodiment, the user may indicate an interest in an attribute or a set of attributes through some user gesture after the initial univariate interestingness sort is performed. In this scenario, following this user gesture, the attributes of the data set are re-sorted, taking into account the new information.

[0053] For the bivariate embodiment, the entropy computation can be based on mutual information. For two attributes X and Y , the mutual information I can be expressed as:

$$I(X;Y) = \sum_{y \in Y} \sum_{x \in X} p(x,y) \log \left(\frac{p(x,y)}{p(x)p(y)} \right),$$

where x and y are possible outcomes for attributes X and Y correspondingly; $p(x)$ and $p(y)$ are probabilities for outcomes x and y correspondingly; $p(x, y)$ is the joint probability for outcomes x and y occurring together (in the same row of data), and the double summation occurs over all possible outcomes.

[0054] As an example for the bivariate embodiment, assume that the data has four attributes: x_1, x_2, x_3, x_4 . Attributes may be discrete or numerical, in which case they can be discretized via binning. For each attribute, the relative entropy (mutual information) is computed relative to the chosen attribute. If x_1 denotes the chosen attribute, then for every other attribute x_2, x_3, x_4 , embodiments can compute mutual information with respect to x_1 . The three attributes can then be sorted according to the computed values of mutual information in respect to the selected column x_1 . For example, if the values of mutual information are $I(x_1, x_2) = 0.4$; $I(x_1, x_3) = 0.6$; $I(x_1, x_4) = 0.2$, then sorting of the attributes according to this mutual information would result in the following ordering of the attributes: x_3, x_2, x_4 . Other embodiments might combine this sorting with other considerations. For example, if it is known that the attributes of city, state and zip code are related to one another, such attributes could be kept together during the sort process.

[0055] In another embodiment, conditional mutual information can be used, where the conditional mutual information in one example is the expected value of the mutual information of two random variables given the value of a third.

[0056] Fig. 6 is a flow diagram of the functionality of guided data exploration module 16 and other elements in accordance with one embodiment of the present invention. Fig. 6 illustrates a bivariate embodiment. The embodiment of Fig. 6 has similar functionality of the embodiment of Fig. 2 at 202, 204, 206, 208, 210 and 212.

At 614, a selection of a tile is received. At 616, the attributes are sorted by level of interestingness relative to the selected tile.

[0057] As disclosed, embodiments provide guided analysis for exploration of datasets. Attributes of a selected dataset are sorted by levels of interestingness, and a user, through an iterative process, can provide further sorting.

[0058] Several embodiments are specifically illustrated and/or described herein. However, it will be appreciated that modifications and variations of the disclosed embodiments are covered by the above teachings and within the purview of the appended claims without departing from the spirit and intended scope of the invention.

WHAT IS CLAIMED IS:

1. A method of exploring data, the method comprising:
 - receiving the data from a database;
 - indexing the data in a server;
 - displaying one or more selectable datasets from the indexed data, each of the datasets comprising one or more attributes;
 - sorting the one or more attributes by level of interestingness using entropy for the attributes; and
 - displaying the sorted attributes.
2. The method of claim 1, wherein at least some of the data in the database is unstructured data.
3. The method of claim 1, wherein the sorting comprises:
 - applying a low-end cutoff to a determined entropy range and applying a high-end cutoff to the determined entropy range, and filtering the entropy based on the low-end cutoff and the high-end cut-off; and
 - mapping low and high values of the filtered entropy to high levels of interestingness, and mapping the intermediate values to low levels of interestingness.
4. The method of claim 1, wherein the displaying the sorted attributes comprises displaying a tile for each sorted attribute, the tile comprising a visualization of the attribute.

5. The method of claim 1, further comprising:
 - receiving a selection of a refinement from the display of sorted attributes and, in response:
 - applying the selected refinement to the data; and
 - re-sorting the one or more attributes of the refined data.
6. The method of claim 1, further comprising performing metadata discovery on the data received from the database.
7. The method of claim 1, further comprising receiving a selection of one or more of the selectable datasets.
8. The method of claim 1, wherein each of the attributes consists of a database column.
9. A computer readable medium having instructions stored thereon that, when executed by a processor, cause the processor to provide guided data exploration, the providing comprising:
 - receiving the data from a database;
 - indexing the data in a server;
 - displaying one or more selectable datasets from the indexed data, each of the datasets comprising one or more attributes;
 - sorting the one or more attributes by level of interestingness using entropy for

the attributes; and
displaying the sorted attributes.

10. The computer readable medium of claim 9, wherein at least some of the data in the database is unstructured data.

11. The computer readable medium of claim 9, wherein the sorting comprises:

applying a low-end cutoff to a determined entropy range and applying a high-end cutoff to the determined entropy range, and filtering the entropy based on the low-end cutoff and the high-end cut-off; and

mapping low and high values of the filtered entropy to high levels of interestingness, and mapping the intermediate values to low levels of interestingness.

12. The computer readable medium of claim 9, wherein the displaying the sorted attributes comprises displaying a tile for each sorted attribute, the tile comprising a visualization of the attribute.

13. The computer readable medium of claim 9, the providing further comprising receiving a selection of a refinement from the display of sorted attributes and, in response:

applying the selected refinement to the data; and
re-sorting the one or more attributes of the refined data.

14. The computer readable medium of claim 9, the providing further comprising performing metadata discovery on the data received from the database.

15. The computer readable medium of claim 9, the providing further comprising receiving a selection of one or more of the selectable datasets.

16. The computer readable medium of claim 9, wherein each of the attributes consists of a database column.

17. A guided data exploration system comprising:
an indexing module that receives data from a database and indexes the data in a server;
a display module that displays one or more selectable datasets from the indexed data, wherein each selectable dataset comprises one or more attributes;
and
a sorting module that sorts the one or more attributes by level of interestingness using entropy for the attributes;
wherein the display module further displays the sorted attributes.

18. The system of claim 17, wherein the sorting comprises:
applying a low-end cutoff to a determined entropy range and applying a high-end cutoff to the determined entropy range, and filtering the entropy based on the low-end cutoff and the high-end cut-off; and

mapping low and high values of the filtered entropy to high levels of interestingness, and mapping the intermediate values to low levels of interestingness.

19. The system of claim 17, wherein the sorting module further receives a selection of a refinement from the display of sorted attributes and, in response:

applies the selected refinement to the data; and

re-sorts the one or more attributes of the refined data.

20. The method of claim 1, wherein the sorting module further receives a selection of one or more of the selectable datasets.





Fig. 1

**Fig. 2**

ORACLE® Big Data Discovery CATALOG		CREATE NEW Project	
675 DATA SETS PROJECTS		301 Data Set	
REFINE BY		<1> PROJECT'S (78) ▼	
USAGE		2 ☆                                             <img alt="Delete icon" data-bbox="315	






Fig. 4

**Fig. 5**

**Fig. 6**

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2015/051462**A. CLASSIFICATION OF SUBJECT MATTER****G06F 17/30(2006.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G06F 17/30; G06T 11/00; G06F 9/44; G06N 5/02Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: dataset, display, attribute, entropy, interestingness, visualization, sort, select, graph, hadoop, refine, and similar terms.**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2011-0302226 A1 (DANIEL ABADI et al.) 08 December 2011 See paragraphs [0037]–[0040] and [0051]; claims 1–6; and figures 1 and 5a–5c.	1–20
A	US 2009-0112904 A1 (DAVOR CUBRANIC) 30 April 2009 See paragraphs [0028]–[0040] and figures 1–2.	1–20
A	US 2014-0218383 A1 (ORACLE INTERNATIONAL CORPORATION) 07 August 2014 See paragraphs [0040]–[0058] and figure 1.	1–20
A	US 2013-0080373 A1 (BIO-RAD LABORATORIES, INC.) 28 March 2013 See paragraphs [0037]–[0060] and figure 2.	1–20
A	US 2007-0094060 A1 (ERIC APPS et al.) 26 April 2007 See paragraphs [0219]–[0235] and figure 13.	1–20

 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	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"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
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 04 February 2016 (04.02.2016)	Date of mailing of the international search report 04 February 2016 (04.02.2016)
Name and mailing address of the ISA/KR International Application Division Korean Intellectual Property Office 189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea Facsimile No. +82-42-472-7140	Authorized officer NHO, Ji Myong Telephone No. +82-42-481-8528

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/051462

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0302226 A1	08/12/2011	EP 2577512 A2 WO 2011-153239 A2 WO 2011-153239 A3	10/04/2013 08/12/2011 16/05/2013
US 2009-0112904 A1	30/04/2009	CA 2667774 A1 CA 2667774 C CN 101606149 A CN 101606149 B EP 2078266 A2 EP 2078266 A4 JP 2010-508593 A JP 2013-020633 A JP 5612047 B2 US 2008-104053 A1 US 2012-221575 A1 US 2013-304766 A1 US 7493330 B2 US 8195695 B2 US 8504589 B2 US 8832153 B2 WO 2008-055114 A2 WO 2008-055114 A3	08/05/2008 17/02/2015 16/12/2009 04/07/2012 15/07/2009 25/08/2010 18/03/2010 31/01/2013 22/10/2014 01/05/2008 30/08/2012 14/11/2013 17/02/2009 05/06/2012 06/08/2013 09/09/2014 08/05/2008 02/10/2008
US 2014-0218383 A1	07/08/2014	None	
US 2013-0080373 A1	28/03/2013	CN 103988078 A EP 2758781 A1 EP 2758781 A4 JP 2014-528080 A WO 2013-044244 A1	13/08/2014 30/07/2014 27/05/2015 23/10/2014 28/03/2013
US 2007-0094060 A1	26/04/2007	EP 1941432 A1 EP 1941432 A4 WO 2007-048229 A1	09/07/2008 20/04/2011 03/05/2007