

DATA RECORDS SELECTION

CLAIM OF PRIORITY

This application claims priority to U.S. Patent Application Serial No. 61/759,799, filed on February 1, 2013, and to U.S. Patent Application Serial No. 13/827,558, filed March 14, 2013, the entire contents of both of which are hereby incorporated by reference.

Background

[0001] Stored data sets often include data for which various characteristics are not known beforehand. For example, ranges of values of typical values for a data set, relationships between different fields within the data set, or functional dependencies among values in different fields, may be unknown. Data profiling can involve examining a source of a data set in order to determine such characteristics.

Summary

[0002] During development of data processing applications, developers may work outside a production environment and may not have access to production data. To ensure that a data processing application (referred to herein as an "application") will execute appropriately in production with actual data, realistic data may be used during execution and testing of the application. Applications often include rules whose execution depends on the value of one or more variables. These variables may be input variables corresponding to input data, may be derived variables that depend on one or more input variables, etc. A subset of data records can be selected from actual production data to be used for development and testing of the application. These data records are generally selected such that the input data is sufficient for every rule in the application to be executed (e.g., such that complete code coverage in the application is achieved).

[0003] In a general aspect, a computer-implemented method includes accessing a plurality of data records, each data record having a plurality of data fields. The method further includes analyzing values for one or more of the data fields for at least some of the plurality of data records and generating a profile of the plurality of data records based on the analyzing. The method further includes formulating at least one subsetting rule

based on the profile; and selecting a subset of data records from the plurality of data records based on the at least one subsetting rule.

[0004] Embodiments may include one or more of the following.

[0005] Formulating at least one subsetting rule includes identifying a first data field as a target data field based on a cardinality of the first data field. In some cases, the target data field has a set of distinct values in the plurality of data records, and wherein selecting a subset of data records includes selecting data records such that there is at least one data record in the selected subset that has each of the distinct values for the target data field.

[0006] Generating a profile includes classifying values for a first data field in the plurality of data records. Formulating at least one subsetting rule includes identifying the first data field as a target data field based on the classifying. In some cases, the target field has a set of distinct values in the plurality of data records, and wherein selecting a subset of data records includes selecting data records such that there is at least one data record in the selected subset that has each of the distinct values for the target data field.

[0007] Formulating at least one subsetting rule includes identifying a first data field as a first target data field and a second data field as a second target data field. In some cases, selecting a subset of data records includes selecting the subset of data records based on combinations of a first set of distinct values for the first target data field and a second set of distinct values for the second target data field.

[0008] Generating a profile includes identifying a relationship between data records related via values of a first data field. The at least one subsetting rule includes an identification of the relationship. In some cases, selecting a subset of data records includes selecting a first data record; and selecting one or more second data records related to the first data record via the relationship identified in the subsetting rule. In some cases, the relationship between data records includes a relationship between data records in a first set of data records and data records in a second set of data records.

[0009] Generating a profile includes generating a a pseudofield for at least some of the plurality of data records; and populating the pseudofield for each corresponding data record with an accumulated value. The accumulated value for a first data record is determined based on the first data record and at least one other data record related to the

first data record. The first data record and the at least one other data record are related via values of a first data field. In some cases, the method includes determining the accumulated value based on a sum of a value for a second data field of the first data record and values for the second data field for each other related data record.

[0010] The method includes receiving a subsetting rule.

[0011] The method includes providing the selected subset of data records to a data processing application. In some cases, the method includes formulating a second subsetting rule based on results of the data processing application; and selecting a second subset of data records based on based on the second subsetting rule.

[0012] In a general aspect, software stored on a computer-readable medium includes instructions for causing a computing system to access a plurality of data records, each data record having a plurality of data fields. The software includes instructions for causing the computing system to analyze values for one or more of the data fields for at least some of the plurality of data records; and generate a profile of the plurality of data records based on the analyzing. The software also includes instructions for causing the computing system to formulate at least one subsetting rule based on the profile; and select a subset of data records from the plurality of data records based on the at least one subsetting rule.

[0013] In a general aspect, a computing system includes at least one processor configured to access a plurality of data records, each data record having a plurality of data fields. The processor is configured to analyze values for one or more of the data fields for at least some of the plurality of data records and generate a profile of the plurality of data records based on the analyzing. The processor is also configured to formulate at least one subsetting rule based on the profile; and select a subset of data records from the plurality of data records based on the at least one subsetting rule.

[0014] In a general aspect, a computing system includes means for accessing a plurality of data records, each data record having a plurality of data fields. The computing system includes means for analyzing values for one or more of the data fields for at least some of the plurality of data records; and means for generating a profile of the plurality of data records based on the analyzing. The computing system also includes means for formulating at least one subsetting rule based on the profile; and means for selecting a

subset of data records from the plurality of data records based on the at least one subsetting rule.

[0015] In a general aspect, a computer-implemented method includes accessing a plurality of data records, each data record having a plurality of data fields, and selecting a first subset of data records from the plurality of data records. The method includes providing the first subset of data records to a data processing application that implements a plurality of rules and receiving a report indicative of a number of times at least one of the rules was executed by the data processing application. The method includes, based on the report, selecting a second subset of data records from the plurality of data records.

[0016] Embodiments may include one or more of the following.

[0017] The method includes providing the second subset of data records to the data processing application.

[0018] The method includes, based on the report, identifying one or more unexecuted rules that were not executed by the data processing application. Selecting the second subset of data records includes selecting data records based on the identifying.

[0019] The method includes, based on the report, identifying one or more rules that were each executed less than a corresponding maximum number of times. Selecting the second subset of data records includes selecting data records based on the identifying.

[0020] The method includes, based on the report, identifying one or more rules that were each executed more than a corresponding minimum threshold number of times. Selecting the second subset of data records includes selecting data records based on the identifying.

[0021] Selecting a first subset of data records includes selecting the first subset of data records based on a first subsetting rule. In some cases, selecting the first subset of data records based on the first subsetting rule includes selecting the first subset of data records such that at least one data record in the subset has each of a set of distinct values for a target data field. In some cases, selecting the first subset of data records based on the first subsetting rule includes selecting a first data record and selecting one or more second data records related to the first data record via a relationship identified in the first subsetting rule. In some cases, selecting the second subset of data records includes

selecting the second subset of data records based on a second subsetting rule different from the first subsetting rule.

[0022] The report includes data indicative of a value of a variable that triggers execution of one or more rules of the data processing application. The method includes identifying one or more data fields as target data fields based on the variable, wherein the variable depends on a value of the identified one or more data fields.

[0023] The second subset of data records includes the first subset of data records.

[0024] The method includes iteratively selecting a subset of data records and providing the subset of data records to the data processing application until the report indicates that at least a threshold number of rules were executed by the data processing application.

[0025] In a general aspect, software stored on a computer-readable medium includes instructions for causing a computing system to access a plurality of data records, each data record having a plurality of data fields, and to select a first subset of data records from the plurality of data records. The software includes instructions for causing the computing system provide the first subset of data records to a data processing application that implements a plurality of rules and receive a report indicative of a number of times at least one of the rules was executed by the data processing application. the software includes instructions for causing the computing system to, based on the report, select a second subset of data records from the plurality of data records.

[0026] In a general aspect, a computing system includes at least one processor configured to access a plurality of data records, each data record having a plurality of data fields, and select a first subset of data records from the plurality of data records. The processor is configured to provide the first subset of data records to a data processing application that implements a plurality of rules and receive a report indicative of a number of times at least one of the rules was executed by the data processing application. The processor is configured to, based on the report, select a second subset of data records from the plurality of data records.

[0027] In a general aspect, a computing system includes means for accessing a plurality of data records, each data record having a plurality of data fields and means for selecting a first subset of data records from the plurality of data records. The computing

system includes means for providing the first subset of data records to a data processing application that implements a plurality of rules and means for receiving a report indicative of a number of times at least one of the rules was executed by the data processing application. The computing system includes means for, based on the report, selecting a second subset of data records from the plurality of data records.

[0028] The techniques described herein have may have one or more of the following advantages. For instance, a complete set of production data records may be massive, and testing a data processing application with such a large set of records can be slow and impractical. By using only a subset of data records selected to be representative of features of the complete set of data records that are relevant to the operation of the data processing application, thorough and efficient testing can be achieved. Accurate selection of a minimum number of data records for effective testing of the application can be achieved via automated profiling analysis of the complete set of data records and feedback from execution of the data processing application.

[0029] Other features and advantages are apparent from the following description and from the claims.

Brief Description of Drawings

[0030] FIG. 1 is a block diagram of a data processing system.

[0031] FIG. 2A is a small portion of an example set of customer transaction records.

[0032] FIG. 2B is a small portion of an example set of demographic records.

[0033] FIG. 3 is a flow chart of an example process for selecting a subset of data records based on target data fields.

[0034] FIG. 4 is a flow chart of an example process for selecting data records.

[0035] FIG. 5 is a flow chart of another example process for selecting data records.

Detailed Description

[0036] During development of data processing applications, developers may work outside of a production environment and may not have access to actual production data. To ensure that a data processing application will execute appropriately in production with

actual data, realistic data may be used during development and testing of the application. Applications often implement rules whose execution depends on (e.g., is triggered by) the value of one or more variables. These variables may be input variables corresponding to input data, may be derived variables that depend on one or more input variables, etc. For effective testing of the application, input data can be provided that is sufficient to cause every logic rule in the application to be executed (e.g., such that complete code coverage in the application is achieved), such that every logic rule is executed at least a corresponding minimum number of times, and/or such that every logic rule is executed no more than a corresponding maximum number of times.

[0037] A subset of data records to be provided to the application is typically selected from one or more larger sets of data records (e.g., from sets of actual production data). The subset may be selected based on subsetting rules, which may be specified by a user, formulated based on a profiling analysis of the data records, formulated based on feedback from execution of an application, etc. For instance, data records that include data that may cause some or all of the rules of the application under test to be executed may be selected for the subset.

[0038] The selected data records are provided to an application, which executes using the selected data records as input data. The application implements one or more rules, i.e., each rule implemented by the application can be executed by the application when a corresponding condition expression for the rule is satisfied and is not executed by the application if the corresponding condition expression is not satisfied. A rule is specified by a specification that includes at least a condition expression and an execution expression. When the condition expression is satisfied (e.g., the condition expression is evaluated with a true result), the execution expression is evaluated. The condition expression may depend on (e.g., is triggered by) the value of one or more variables, which may be input variables corresponding to input data, may be derived variables that depend on one or more input variables, etc. In some examples, the application executes all of the rules that are triggered. In some examples, the application executes fewer than all of the rules that are triggered, such as some of the rules or only one of the rules (e.g., a first rule that is triggered). Rules are described in more detail at least at col. 5, line 61 – col. 6, line 11 of U.S. Patent No. 8,069,129, filed April 10, 2007, the contents of which are incorporated herein by reference in their entirety.

[0039] After execution, a report can be provided which contains data indicative of the execution of the application (e.g., rules in the application did or did not execute, a number of times each logic rule in the application was executed, or other execution data). Based on the report, additional input data may be identified, e.g., input data that would have caused the unexecuted rules to be executed, input data that would have caused a particular logic rule to be executed a specified number of times, or input data that would have caused another desired execution result. Corrective action can be executed, for example, additional subsetting rules can be formulated and an updated subset of data records can be selected according to those additional subsetting rules. The updated subset of data records may include data records sufficient to cause execution of some or all of the previously unexecuted rules, data records sufficient to cause execution of some or all of the rules a specified number of times, or data records sufficient to cause another desired execution result.

[0040] Referring to Fig. 1, a data processing system 100 includes a record selection subsystem 102 hosted on a server 102a. The record selection subsystem 102 selects data records from one or more sets of data records (e.g., production data records). The selected data records are provided to a data processing application 106, e.g., an application that is undergoing testing or development. In some examples, the application 106 is local to the record selection subsystem 102, for example, hosted on the same server 102a. In some examples, the application 106 may be remote to the record selection subsystem 102, for example, hosted on a remote server 106a that is accessed over one or more networks such as a local or wide area data network 118 (e.g., the Internet).

[0041] The data records are stored in data sources 104 hosted by one or more servers 104a, 104b, 104c, 104d and corresponding storage devices 108a, 108b, 108c, 108d. The data sources 104 may include any of a variety of data sources, such as a database 109, a spreadsheet file 110, a text file 112, a native format file 114 used by a mainframe, or another type of data source. One or more of the data sources may be local to the record selection subsystem 102, for example, hosted on the same computer system (e.g., server 102a). One or more of the data sources may be remote to the record selection subsystem 102, for example, hosted on a remote computer (e.g., servers 104a, 104b, 104c, 104d) that is accessed over the network 118, multiple networks, etc.

[0042] The data records stored in data sources 104 include one or more sets of data records. For example, the data records may include customer transaction records, customer demographic records, financial transaction records, telecommunications data, or other types of data records. Each data record has one or more data fields, and each data field has a particular value (or lack thereof) for each data record, such as a numerical value, an alphanumeric value, a null value, etc. For instance, in a set of customer transaction records, each record may have data fields storing a customer identifier, a purchase price, and a transaction type, among other data.

[0043] A subsetting module 120 in the record selection subsystem 102 may provide various operations such as selecting a subset of data records from one or more sets of data records stored in one or more of the data sources 104 according to one or more subsetting rules. A subsetting rule is a rule, executable by a computer, by which a subset of data records from the one or more sets of data records is to be selected. The subsetting rules may be formulated by the subsetting module 120 based on an analysis of a profile of the one or more sets of data records generated by a profiling module 126. The subsetting rules may also be formulated by the subsetting module 120 based on an analysis of the results of the execution of the application provided by a coverage analysis module 128 (e.g., based on the report). The subsetting rules may be specified by a user via a user interface 124, e.g., based on the user's understanding of the data records and/or the application 106 under test. The subsetting rules may also be read from a storage medium, such as a hard disk, or received via a network, such as the Internet.

[0044] A wide variety of subsetting rules are possible and may be applied alone or in combination. Subsetting rules can be deterministic (e.g., the rule may specify that all records matching a particular criterion are to be selected) or non-deterministic (e.g., the rule may specify that, of all the records that match a particular criterion, two of those records are to be selected at random).

[0045] In some examples, a subsetting rule designates one or more target data fields and specifies that each distinct value or value classification for the target data fields is to be included in at least one of the data records of the selected subset of data records. The subsetting module 120 identifies each distinct value for the target data fields in the one or more sets of data records and selects data records to satisfy the subsetting rule. For instance, a *state* data field, which has a distinct value for each of the fifty states, and a

gender data field, which has two distinct values may be identified as target data fields. Data records are selected for the subset such that each of the fifty values for state and each of the two values for gender are included in at least one data record in the subset.

[0046] In some examples, a subsetting rule specifies a type of relationship among data records within the same set of data records or between different sets of data records. The subsetting module 120 selects data records based on their relationship with other data records selected for the subset. For instance, data records that share a common value for a customer identifier (*cust_id*) data field may be selected for the subset. Other examples of subsetting rules are also possible, such as filtering. In some examples, a combination of subsetting rules can be used to select data records for the subset.

[0047] In some examples, subsetting rules are provided by a user, such as a data analyst or an application developer. For instance, a user may identify target fields, specify relationships among data records, or otherwise indicate a subsetting rule.

[0048] In some examples, subsetting rules are formulated by the subsetting module 120 based on an analysis of a profile of the data records automatically generated by the profiling module 126. The profiling module 126 may access the one or more sets of data records and generate a profile of the data records by analyzing individual data records of a single data set and/or analyzing relationships among data fields within a set of data records and/or across different sets of data records.

[0049] A profile of a set of data records is a summary, e.g., on a field-by-field basis, of data in the set of data records. A profile can include information characterizing the data in the set of data records, such as a cardinality of one or more of the data fields in the data records, a classification of values in one or more of the data fields, relationships among data fields in individual data records, relationships among data records, or other information characterizing the data in the set of data records. A profile of a set of data records may also include information characterizing a pseudofield, which is a data field generated by the profiling module 126 and populated with values determined by manipulation of values for one or more data fields in related data records.

[0050] Based on the generated profile of the data records, the subsetting module 120 may identify features of the data records that may be relevant to the selection of a subset of data records that achieves good code coverage for the application 106. For instance,

based on the profile of the data records, the subsetting module 120 may identify one or more data fields or combinations of data fields that are likely to relate to the input variables and derived variables of the application. In some cases, subsetting rules may also be formulated based on input received from a user or from a computer storage medium and/or based on results of execution of the application 106 (e.g., based on input received from a coverage analysis module 128).

[0051] The subsetting module 120 may execute operations for one or more types of analysis to specify subsetting rules. The subsetting module 120 may specify one or more subsetting rules based on an analysis of the data fields within individual data records, e.g., by determining which data fields are likely to relate to variables in the application 106. In some examples, the subsetting module 120 identifies a target data field based on the cardinality of the target data field as indicated in the profile (i.e., the number of distinct values or classifications of values for the data field across all of the data records of a set). For instance, a *gender* data field (with a cardinality of two) may be identified as a target data field while a *phone_number* data field (with a cardinality on the order of the total number of data records) is not likely to be identified as a target data field. In some examples, subsetting module 120 identifies as a target data field a pseudofield populated with data resulting from a manipulation of data in one or more data fields. For instance, data in an *income* data field may be classified into categories (e.g., high, medium, or low), and a pseudofield populated with the classifications of the *income* data field (*inc_range*) may be identified as a target data field. In some examples, the subsetting module 120 identifies a target data field based on relationships between the target data field and one or more other data fields within the same record as indicated in the profile. For instance, the profile may indicate that the data fields *state* and *ZIP* are not independent; based on this dependence, the subsetting module 120 may consider only one of those data fields as a possible target data field. The subsetting module 120 may also specify one or more subsetting rules based on an analysis of relationships among different data records within a set of data records and/or across different sets of data records as indicated in the profile. For instance, the profile may indicate that data records may be linked via a common value of a data field (e.g., the value of a *cust_id* data field). Other analyses of the data records are also possible.

[0052] Once a subset of data records is selected by the subsetting module 120, data indicative of the selected subset of data records are provided to the application 106 under

test. For instance, identifiers for the selected subset of data records and an address for the data records may be provided to the application 106. A file containing the selected subset of data records may also be provided to the application 106.

[0053] The data processing application 106 is executed using the subset of data records as input data. After execution, a report is provided to a coverage analysis module 128 in the record selection subsystem 102. The report may also be provided to the user 122. The report contains data indicative of the execution of the application (e.g., rules in the application that did or did not execute, a number of times each logic rule in the application was executed, or other execution data). In some examples, the report directly identifies those rules that did or did not execute. The report may also contain additional information about the execution of the application 106, such as a number of times each logic rule was executed, a value of each variable of the application during execution, or other information.

[0054] For each logic rule in the application that did not execute, the coverage analysis module 128 identifies one or more variables of the application 106 that relate to that logic rule. The coverage analysis module 128 may identify variables based on data included in the report (e.g., data indicative of the flow of data through the application 106), based on preloaded information about the application, etc. In some cases, the coverage analysis module 128 also identifies a value or range of values for each variable that would have caused the logic rule to execute. The input data fields and values or ranges of values that correspond to the variables are identified and used to specify additional subsetting rules in subsequent selection of an updated subset of data records by the subsetting module 120.

[0055] For example, if an identified variable is an input variable of the application that directly corresponds to one of the data fields of the data records, the coverage analysis module 128 identifies the corresponding data field and a value or range of values for the data field. For example, if a logic rule in the application 106 executes when a variable *x* is greater than 10, and the variable *x* corresponds to an input data field *txn_amt* containing data about the amount of a customer transaction, the coverage analysis module determines that the input data should include at least one data record for which *txn_amt* > 10. This determination (e.g., *txn_amt* > 10) is provided to the subsetting module 120, which specifies an additional subsetting rule such that subsequent subsets of data records

provided to the application 106 will include data sufficient to cause the $x > 10$ logic rule to execute.

[0056] For example, if an identified variable is not an input variable (i.e., the identified variable does not directly correspond to one of the data fields of the data records), a data lineage submodule 130 in the coverage analysis module 128 tracks the derivation of the variable through the logic of the application 106 to identify the input variable(s) from which the identified variable is derived. The coverage analysis module 128 then identifies the corresponding data field(s) and values or ranges of values for the data field. For example, if a logic rule in the application 106 executes when the value of a variable y is 2, the data lineage submodule 130 may determine that y is derived via logical steps in the application from a logical combination of three input variables that correspond to input data fields *gender*, *inc_range*, and *state*. By following the logical derivation of the variable y , the values of the data fields *gender*, *inc_range*, and *state* that cause $y = 2$ can be determined. For instance, the logic rule $y = 2$ may be satisfied when *gender* = F, *inc_range* = high, and *state* = ME, NH, VT, MA, RI, or CT. The determination is provided to the subsetting module 120, which specifies an additional subsetting rule such that subsequent subsets of data records provided to the application 106 will include data sufficient to cause the $y = 2$ logic rule to execute. As another example, a logic rule may execute when the values of two variables have a particular relationship, such as when values of variables corresponding to the data fields *firstname* and *lastname* are equal.

[0057] In some examples, the results of the coverage analysis are also provided to the user 122. The user may provide additional subsetting rules to the subsetting module 120 or may modify previously provided subsetting rules. The user may also provide additional input to the profiling module 126 to modify the input previously provided to the profiling module.

[0058] In some examples, even the complete set(s) of data records does not include data sufficient to satisfy a logic rule in the application 106. For instance, the application 106 may include a logic rule that executes only when the value for the data field *income* is greater than five million dollars. If no data record exists in the set for which *income* > \$5,000,000, then no subset of data records will cause that logic rule to execute. To identify such deficiencies in the complete data set, in some examples, the application may

be executed one or more times using all of the data records as input. The resulting report identifies rules that cannot be covered regardless of the subset of data records that are selected for input.

[0059] The operation of the subsetting module 120 and the profiling module 126 will be described with reference to example sets 200, 252 of data records shown in Figs. 2A and 2B. Fig. 2A is an example of a small portion of a set 200 of customer transaction records. Each customer transaction record 202 has several data fields 204, including, e.g., a customer identifier (*cust_id*) 204a, a transaction type (*txn_type*) 204b, a transaction amount (*txn_amt*) 204c, a transaction date (*date*) 204d, and a store identifier (*store_id*) 204e. Other data fields may also be included. Fig. 2B is an example of a small portion of a set 250 of demographic records. Each demographic record 252 has several data fields 254, including, e.g., a customer identifier (*cust_id*) 254a, a customer address (*address*, *state*, *ZIP*) 254b, 254c, 254d, a customer income (*income*) 2054e, and a customer gender (*gender*) 254f. Other data fields may also be included. The operations of the profiling module 126 and the subsetting module 120 are not restricted to these example data sets and apply similarly to other types of data sets.

[0060] The subsetting module 120 may select a subset of data records according to one or more types of subsetting rules. Some example subsetting rules are as follows:

[0061] **Filtering.** In some examples, the subsetting module 120 selects a subset of data records according to a filter. For instance, the filter may specify that all data records having a particular value for a given data field are to be selected. For example, the filter may specify that all demographic records from the set 250 having *state* (data field 254c) = "MA" are to be selected for the subset. The filter may be specified by the user, the profiling module 126, and/or the coverage analysis module 128.

[0062] In some examples, the subsetting module 120 selects a subset of data records according to a rules-based filter in which data records are eliminated based on a value of a given data field. For example, the filter may specify that data records for which *store_id* (data field 204e) = "online" are to be eliminated from the subset. The rules-based filter may be specified by the user 122, the profiling module 126, and/or the coverage analysis module 128.

[0063] Target Data Fields. In some examples, the subsetting module 120 selects a subset of data records based on one or more target data fields. Target data fields are e.g., data fields that are likely to be related to variables of the application. For example, if a particular application that operates on customer transaction records tracks transaction type (i.e., purchase or return) by store location, the developer of the application may identify the data fields *txn_type* (data field 204c) and *store_id* (data field 204e) as target data fields. In some cases, the subsetting module 120 may identify a target data field based on characteristics of the data field, such as a cardinality of the data field, as indicated in a profile of the data records. In some cases, the coverage analysis module 128 may identify a target data field based on a relationship between a variable of the application and a data field. A data field of low cardinality (e.g., data fields with cardinality less than a threshold cardinality) may be identified as a target data field, even if the profiling module 126 has little or no other information about the content of the data field and how that content may relate to an application. The threshold cardinality may be specified by a user or may be determined automatically by the profiling module. For instance, based on a profile of the set 350 of demographic records, the data field *state* may be identified as a target data field provided the threshold cardinality is set to at least fifty.

[0064] Fig. 3 is a flow chart of an example process for selecting a subset of data records based on target data fields. One or more target data fields are identified (300), e.g., based on information included in a profile of the data records, information from a user, information from the coverage analysis module 128, etc. A set of distinct values for each target data field in the set of records is identified (302). Data records are selected for the subset (304) such that each distinct value for each target data field is included in at least one data record in the subset. In an example, the *state* data field and the *gender* data field are identified as target data fields for the set 250 of demographic records. The set 250 of data records is analyzed to identify fifty distinct values for *state* and two distinct values for *gender*. Data records are selected such that each of the fifty values for *state* and each of the two values for *gender* are included in at least one data record in the subset. In some examples, the subsetting rule may specify a number of times that each distinct value for each target data field is included in the subset (e.g., one time, ten times, fifty times, etc.).

[0065] Subsetting based on target data fields does not necessarily imply that every combination of each value of each data field is represented in the subset. For instance, a subset of data records in which each of the fifty values for *state* and each of the two values for *gender* are included may contain only fifty data records. In some examples, the target data field is a constructed field, such as a pseudofield (e.g., constructed by the profiling module, as described below), and depends on one or more data fields within the same record or across different records.

[0066] **Data Classification.** In some examples, a subset of data records is selected based on a classification of data in one or more target data fields of a data record. For instance, a subsetting rule may identify a target data field and specify distinct ranges ("bins") of values by which the values of the target data field can be classified. Data records are selected for the subset based on the bins of the target data field rather than the precise values of the target data field. In an example, the data field *income* in the set 250 of demographic records is identified as a target data field. Three bins are designated: "low" (*income* < \$50,000), "medium" (*income* between \$50,000 and \$150,000), and "high" (*income* > \$150,000). The value of the *income* data field in each data record considered by the subsetting module 120 for inclusion in the subset is classified as low, medium, or high; and data records are selected such that each of the three bins of *income* is included in at least one data record in the subset. In some examples, values of a data field are classified (e.g., by the profiling module) and a pseudofield for each data record is populated with the corresponding classified value (e.g., the data field *inc_range* 256). In these examples, the pseudofield is treated as a target data field and data records are selected such that each distinct value for the pseudofield is included in at least one data record in the subset. The data field to be classified, the number of bins, and/or the range of values for each bin may be specified by the user 122 or identified automatically by the profiling module 126 and/or the coverage analysis module 128.

[0067] **Combinatorics.** In some examples, a subset of data records is selected according to a combinatorics rule, which may specify combinations of two or more other subsetting rules. For instance, a combinatorics rule may identify two target data fields and specify that all possible combinations of all values of each of the two target data fields be included in at least one data record in the subset. An example combinatorics rule may identify the data fields *inc_range* and *gender* as target data fields and specify that all possible combinations of those two data fields be included in the subset. The subset

satisfying this combinatorics rule would include six data records (i.e., low + female, low + male, middle + female, middle + male, high + female, high + male). In contrast, without the combinatorics rule, the specification of *inc_range* and *gender* as target data fields could be satisfied with as few as three records (e.g., low + female, middle + male, high + female). In some examples, a subsetting rule may specify a combinatoric combination of two or more target data fields and one or more other target data fields outside of the combinatoric combination. For instance, a subsetting rule may specify *inc_range* and *gender* as target data fields to be taken in combinatoric combination, and may also specify *state* as a target data field outside of that combination. More complex combinations are also possible. The target data fields and the specific type of combination may be specified by the user 122 or identified automatically by the profiling module 126 and/or the coverage analysis module 128.

[0068] Relationships Among Data Records. In some examples, a subset of data records is selected according to a relationship among data records within a set of data records or across different sets of data records. A subsetting rule may specify a join key such that if one data record is selected for the subset, other data records that are related to that data record via the join key are also selected for the subset. For instance, a subsetting rule may identify the data field *cust_id* as a join key relating data records within the set 200 of customer transaction records and between the set 200 and the set 250 of demographic records. For each data record from either set that is selected for the subset (e.g., according to another subsetting rule), other data records that share the same value for *cust_id* as the selected data record are also selected for the subset. By selecting data records according to relationships, the subset will contain, e.g., data records for all transactions for a particular customer as well as demographic records for that customer. The relationship may be specified by the user 122 or identified automatically by the profiling module 126 and/or the coverage analysis module 128.

[0069] In some examples, the relationship among data records may be based on one or more characteristics of the data records. For instance, a data record of interest may be identified (e.g., a data record corresponding to a fraudulent credit card transaction). A corresponding subsetting rule may then specify that the subset is to include fifty other data records having characteristics similar to the identified data record of interest, e.g., to help identify other instances of fraud in the data records.

[0070] Other subsetting rules may also be specified. For instance, data record counts may be specified (e.g., the subset is to include at least 100 records for which *txn_type* = "purchase"). Statistical parameters may be specified (e.g., the subset is to include all data records for which *txn_type* = "purchase" and 15% of data records for which *txn_type* = "return"). Numerical parameters may be specified (e.g., the subset is to include at least a specified number of data records per million data records in the set of data records). These subsetting rules may be specified by the user 122 and/or formulated by the subsetting module 120 based on an analysis of the profile (generated by the profiling module 126) and/or the results of an analysis of the execution (provided by the coverage analysis module 128).

[0071] In some examples, multiple subsetting rules may be applied to a set of data records. In some cases, application of these multiple subsetting rules may result in some data records being selected multiple times for the subset. A deduplication rule may be applied to the selected data records to remove any data records that appear more than once in the subset.

[0072] In some examples, subsetting rules are formulated based on an analysis of the profile generated by the profiling module 126. The profiling module 126 may analyze the data records with no input from outside sources or with input from the user 122 and/or the coverage analysis module 128. Some examples of profiling analyses are as follows:

[0073] **Cardinality.** In some examples, the profiling module 126 identifies the cardinality of the data field (i.e., the number of distinct values for the data field across all of the data records of a set). For instance, when profiling the set 300 of customer transaction records, the profiling module may identify *txn_type* as a data field of low cardinality (only two distinct values in all the data records of the set 300). When profiling the set 350 of demographic records, the data field *state* may be identified as a data field with cardinality fifty, provided the threshold cardinality is set to at least fifty. The cardinality of some or all of the data fields, may be used by the subsetting module 120 to specify a subsetting rule.

[0074] **Classification.** In some examples, the profiling module 126 classifies data in a data field. For instance, the profiling module may identify distinct ranges ("bins") of values by which the values of a high cardinality data field can be classified. As classified, the data field has a lower cardinality and thus may be identified as a target data field as

described above. In some cases, the profiling module classifies the value for the data field of each record as it analyzes the record but does not store the classification. In some cases, the profiling module generates a pseudofield for each record in which the bin corresponding to value of the data field is stored. As an example, the data field *income* in the set 350 of demographic records is high cardinality. The profiling module classifies the *income* value for each record into one of three bins (high, medium, or low) and generates a pseudofield *inc_range* 356 to store the categorized data. The pseudofield 356 has cardinality three and as such may be identified as a target data field by the subsetting module 120 where the high cardinality data field *income* may not have been identified as a target data field. In some examples, the profiling module recognizes that a high cardinality data field can be classified automatically. In some examples, a user identifies a data field for classification and may also specify the number of bins and the ranges of values that fall within each bin. In some examples, the user specifies characteristics of data fields to be classified without identifying a specific data field (e.g., the user may specify that any data field having numeric values and having cardinality between 10 and 100 is to be categorized into quartiles).

[0075] Relationships Among Data Fields. In some examples, the profiling module 126 determines relationships among data fields within a single data record. For instance, if a first data field in a data record depends on a second data field within each data record, then only one of the first data field and the second data field needs to be considered as a target data field. For example, the data field *state* and the data field *ZIP* are related (i.e., the value for *ZIP* depends on the value for *state*). Based on an indication of such a relationship in the profile, the subsetting module 120 may consider only one of the two related data fields as a potential target data field. More complex relationships among data fields can also be identified and used by the subsetting module 120 in identifying target data fields. The profiling module may be guided by user input, e.g., by user specification of data fields that are likely to be related.

[0076] Relationships Among Data Records. In some examples, the profiling module 126 determines relationships among different data records within a set of data records or across different sets of data records. For instance, the profiling module may recognize that some data records within a set are linked via a common value of a data field. For example, the set 300 of customer transaction records may include multiple data records corresponding to transactions by the same customer. These data records are

linked via a common value for *cust_id* (i.e., a join key). The profiling module may also recognize that a first data record within a first set is related via a common value of a data field to a second data record within a second set. For example, a data record in the set 300 of customer transaction records may be linked to a data record in the set 350 of demographic records via the data field *cust_id* (i.e., a transaction record for a particular customer can be linked to the demographic record for that customer). The profiling module may be guided by user input, e.g., by user specification of data fields that are likely to link data records. The profiling module may also be guided to identify join keys or other relationships via an analysis of the schema of a relational database associated with the set(s) of data records. In some examples, the profiling module 126 determines relationships among data records and presents the relationships to the user, who may then use the information about relationships to specify subsetting rules to the subsetting module 120.

[0077] Based on an indication of such a relationship among data records in the profile, the subsetting module 120 may specify the join key as part of a subsetting rule. Under such a subsetting rule, if one data record is selected for the subset, other data records related to that data record via the join key are also selected for the subset (e.g., if one data record having a given *cust_id* is selected for the subset, other data records having the same *cust_id* are also selected).

[0078] **Pseudofields.** In some examples, the profiling module 126 generates a new pseudofield with values determined by manipulation of values for one or more data fields in related data records and identifies the pseudofield as a target data field. The value for the pseudofield may be a combination of values for one or more data fields in data records that are related via a join key. For instance, the value for the pseudofield may be an accumulated value, e.g., an accumulation, such as a sum, count, or other accumulation, of all values for a first data field for data records that are related via a common value of a second data field. The value for the pseudofield may also be a classification of the accumulated value. For example, to address logic in the application that performs an action depending on a total transaction amount for a given customer, a pseudofield *total_amt* 306 is generated in the set 300 of customer transaction records. The value for the pseudofield *total_amt* for a data record having a given *cust_id* value is determined by summing the values for the *txn_amt* field for all data records having that *cust_id* value

and classifying the sum into one of three bins (high, medium, or low). The pseudofield may then be identified by the subsetting module as a target data field.

[0079] Referring to Fig. 4, in an example process, a plurality of data records are accessed (400). Each data record has a plurality of data fields. Values for one or more of the data fields for at least some of the plurality of data records are analyzed (402). A profile of the plurality of data records is generated (404) based on the analyzing. A profile of a plurality of data records includes information characterizing the data in the set of data records. At least one subsetting rule is formulated based on the profile (406). A subsetting rule is a specification of a rule by which a subset of data records from the plurality of data records is to be selected. A subset of data records is selected (406) based on the at least one subsetting rule. For instance, the subset of data records may be selected based on values of a target data field and/or based on a relationship between data records that are related via values of a data field.

[0080] Referring to Fig. 5, in another example process, a plurality of data records are accessed (500). Each data record has a plurality of data fields. A first subset of data records is selected from the plurality of data records (502). The first subset of data records is provided to a data processing application (504), such as an application under test. The application implements a variety of rules. Rules in a data processing application are executable portions of the application whose execution depends on (e.g., is triggered by) the value of one or more variables. A report is received (506) indicative of a number of times at least one of the rules was executed by the data processing application. Based on the report, a second subset of data records is selected (508) from the plurality of data records. The second subset of data records is provided to the data processing application (510). For instance, the second subset is selected such that the previously non-executed rules may be executed, or such that certain rules may be executed.

[0081] In some examples, new data records can be generated based on the profiling analysis conducted by the profiling module 126. For instance, the profiling analysis reveals relationships among data fields within and among data records and ranges of possible values for the data fields in an existing set of data records. A new data record can be constructed in which at least some of the data fields are populated with values calculated or determined from the information about existing data records. Test data generation may be used, for instance, when there are no data records in the source data

set that will cause a particular logic rule in an application to execute, e.g., a logic rule that requires *income* > \$10,000,000; or a logic rule that requires a complex combination of particular values of multiple data fields where all of the required values are not represented in the set of data records. Test data generation may also be used to generate a new data set whose profile matches the profile of an original data set. For instance, a new data set may be generated by randomizing the data of the original data set to preserve the privacy of the original data records.

[0082] In some examples, the approaches described above are implemented in an execution environment which may be hosted on one or more general-purpose computers under the control of a suitable operating system, such as the UNIX operating system. For example, the execution environment can include a multiple-node parallel computing environment including a configuration of computer systems using multiple central processing units (CPUs), either local (e.g., multiprocessor systems such as SMP computers), or locally distributed (e.g., multiple processors coupled as clusters or MPPs), or remote, or remotely distributed (e.g., multiple processors coupled via a local area network (LAN) and/or wide-area network (WAN)), or any combination thereof.

[0083] In some cases, the approaches described above are implemented by a system for developing applications as dataflow graphs that include vertices (representing components or datasets) connected by directed links (representing flows of work elements) between the vertices. For example, such an environment is described in more detail in U.S. Publication No. 2007/0011668, entitled “Managing Parameters for Graph-Based Applications,” incorporated herein by reference. A system for executing such graph-based computations is described in U.S. Patent 5,566,072, EXECUTING COMPUTATIONS EXPRESSED AS GRAPHS, incorporated herein by reference. Dataflow graphs made in accordance with this system provide methods for getting information into and out of individual processes represented by graph components, for moving information between the processes, and for defining a running order for the processes. This system includes algorithms that choose interprocess communication methods (for example, communication paths according to the links of the graph can use TCP/IP or UNIX domain sockets, or use shared memory to pass data between the processes).

[0084] The approaches described above can be implemented using software for execution on a computer. For instance, the software forms procedures in one or more computer programs that execute on one or more programmed or programmable computer systems (which may be of various architectures such as distributed, client/server, or grid) each including at least one processor, at least one data storage system (including volatile and non-volatile memory and/or storage elements), at least one input device or port, and at least one output device or port. The software may form one or more modules of a larger program, for example, that provides other services related to the design and configuration of dataflow graphs. The nodes and elements of the graph can be implemented as data structures stored in a computer readable medium or other organized data conforming to a data model stored in a data repository.

[0085] The software may be provided on a storage medium, such as a CD-ROM, readable by a general or special purpose programmable computer, or delivered (encoded in a propagated signal) over a communication medium of a network to a storage medium of the computer where it is executed. All of the functions may be performed on a special purpose computer, or using special-purpose hardware, such as coprocessors. The software may be implemented in a distributed manner in which different parts of the computation specified by the software are performed by different computers. Each such computer program is preferably stored on or downloaded to a storage media or device (e.g., solid state memory or media, or magnetic or optical media) readable by a general or special purpose programmable computer, for configuring and operating the computer when the storage media or device is read by the computer system to perform the procedures described herein. The inventive system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer system to operate in a specific and predefined manner to perform the functions described herein.

[0086] A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, some of the steps described above may be order independent, and thus can be performed in an order different from that described.

[0087] It is to be understood that the foregoing description is intended to illustrate and not to limit the scope of the invention, which is defined by the scope of the appended claims. For example, a number of the function steps described above may be performed in a different order without substantially affecting overall processing. Other embodiments are within the scope of the following claims.

CLAIMS

1. A computer-implemented method for selecting data records to cause execution of a processing rule during testing of a data processing application, the method including:
 - obtaining a first set of data records;
 - processing the first set of data records using a data processing application that includes a processing rule, wherein a processing rule operates on at least one input value and generates at least one output value, and wherein whether the processing rule is executed by the data processing application during processing of a particular data record depends directly or indirectly on a value in each of one or more data fields of the particular data record;
 - receiving execution information indicative of a number of times the processing rule was executed in connection with processing of the first set of data records;
 - obtaining a second set of data records;
 - analyzing values in one or more data fields of each of the data records in the second set, the analyzing including generating a profile of each of one or more of the data fields for the second set of data records, the profile of a data field characterizing the values in the data field;
 - obtaining a subsetting rule based on (i) the generated profile and (ii) the execution information indicative of the number of times the processing rule was executed in connection with processing the first set of data records, the subsetting rule including an identification of a particular one of the data fields of the data records in the second set;
 - selecting a subset of data records from the second set of data records according to the subsetting rule, the selecting of the subset of data records being based on values in the particular one of the data fields; and
 - processing the selected subset of data records using the data processing application.
2. The method of claim 1, wherein obtaining the subsetting rule includes formulating the subsetting rule, including identifying the particular one of the data fields based on a cardinality of the particular identified one of the data fields.
3. The method of claim 2, wherein the particular one of the data fields has a set of distinct values in the plurality of data records, and wherein selecting a subset of data records includes selecting data records such that there is at least one data record in the selected subset that has each of the distinct values for the particular one of the data fields.
4. The method of claim 1, wherein generating a profile includes classifying values for a first data field of the data records in the second set of data records; and

wherein obtaining the subsetting rule includes formulating the subsetting rule, including identifying the first data field as the particular one of the data fields based on the classifying.

5. The method of claim 4, wherein the particular one of the data fields has a set of distinct values of the data records in the second set of data records, and wherein selecting a subset of data records includes selecting data records such that there is at least one data record in the selected subset that has each of the distinct values for the particular one of the data fields.

6. The method of claim 1, wherein the subsetting rule identifies a first data field as a first particular data field and a second data field as a second particular data field, and wherein selecting a subset of data records includes selecting the subset of data records based on combinations of a first set of distinct values for the first particular data field and a second set of distinct values for the second particular data field.

7. The method of claim 1, wherein generating a profile includes identifying a relationship between data records of the second set of data records related via values of a first data field; and wherein the at least one subsetting rule includes an identification of the relationship.

8. The method of claim 7, wherein selecting a subset of data records includes:
selecting a first data record; and
selecting one or more second data records related to the first data record via the relationship identified in the subsetting rule.

9. The method of claim 7, wherein the relationship between data records includes a relationship between data records in the second set of data records and data records in a third set of data records.

10. The method of claim 1, wherein generating a profile includes:
generating a pseudofield for at least some of the data records in the second set of data records;
and
populating the pseudofield for each corresponding data record with an accumulated value,
wherein the accumulated value for a first data record is determined based on the first data record and at least one other data record related to the first data record,
wherein the first data record and the at least one other data record are related via values of a first data field.

11. The method of claim 10, including determining the accumulated value based on a sum of a value for a second data field of the first data record and values for the second data field for each other

related data record.

12. The method of claim 1, including:
 - formulating a second subsetting rule based on results of processing the selected subset of data records by the data processing application; and
 - selecting a second subset of data records based on the second subsetting rule.
13. A non-transitory computer-readable medium storing instructions for causing a computing system to execute the method of any one of claims 1 to 12.
14. A computing system for selecting data records to cause execution of a processing rule during testing of a data processing application, the system including:
 - at least one processor configured to:
 - obtain a first set of data records;
 - process the first set of data records using a data processing application that includes a processing rule, wherein a processing rule operates on at least one input value and generates at least one output value, and wherein whether the processing rule is executed by the data processing application during processing of a particular data record depends directly or indirectly on a value in each of one or more data fields of the particular data record;
 - receive execution information indicative of a number of times the processing rule was executed in connection with processing of the first set of data records;
 - obtain a second set of data records;
 - analyze values in one or more data fields of each of the data records in the second set, the analyzing including generating a profile of each of one or more of the data fields for the second set of data records, the profile of a data field characterizing the values in the data field;
 - obtain a subsetting rule based on (i) the generated profiles and (ii) the execution information indicative of the number of times the processing rule was executed in connection with processing the first set of data records, the subsetting rule including an identification of a particular one of the data fields of the data records in the second set;
 - select a subset of data records from the second set of data records according to the subsetting rule, the selecting of the subset of data records being based on values in the particular one of the data fields; and
 - processing the selected subset of data records using the data processing application.
15. The computing system of claim 14, wherein obtaining the subsetting rule includes formulating the subsetting rule, including identifying the particular one of the data fields based on a cardinality of the particular identified one of the data fields.

16. The computing system of claim 15, wherein the particular one of the data fields has a set of distinct values in the plurality of data records, and wherein selecting a subset of data records includes selecting data records such that there is at least one data record in the selected subset that has each of the distinct values for the particular one of the data fields.

17. The computing system of claim 14, wherein generating a profile includes classifying values for a first data field of the data records in the second set of data records; and

wherein obtaining the subsetting rule includes formulating the subsetting rule, including identifying the first data field as the particular one of the data fields based on the classifying.

18. The computing system of claim 17, wherein the particular one of the data fields has a set of distinct values of the data records in the second set of data records, and wherein selecting a subset of data records includes selecting data records such that there is at least one data record in the selected subset that has each of the distinct values for the particular one of the data fields.

19. The computing system of claim 14, wherein the subsetting rule identifies a first data field as a first particular data field and a second data field as a second particular data field, and wherein selecting a subset of data records includes selecting the subset of data records based on combinations of a first set of distinct values for the first particular data field and a second set of distinct values for the second particular data field.

20. The computing system of claim 14, wherein generating a profile includes identifying a relationship between data records of the second set of data records related via values of a first data field; and

wherein the at least one subsetting rule includes an identification of the relationship.

21. The computing system of claim 20, wherein selecting a subset of data records includes:

selecting a first data record; and

selecting one or more second data records related to the first data record via the relationship identified in the subsetting rule.

22. The computing system of claim 14, wherein generating a profile includes:

generating a pseudofield for at least some of the data records in the second set of data records;

and

populating the pseudofield for each corresponding data record with an accumulated value, wherein the accumulated value for a first data record is determined based on the first data record and at least one other data record related to the first data record,

wherein the first data record and the at least one other data record are related via values of a first data field.

23. The computing system of claim 22, including determining the accumulated value based on a sum of a value for a second data field of the first data record and values for the second data field for each other related data record.

24. The computing system of claim 14, including:
formulating a second subsetting rule based on results of processing the selected subset of data records by the data processing application; and
selecting a second subset of data records based on the second subsetting rule.

25. A computing system for selecting data records to cause execution of a processing rule during testing of a data processing application, the computing system including:

means for obtaining a first set of data records;

means for processing the first set of data records using a data processing application that includes a processing rule, wherein a processing rule operates on at least one input value and generates at least one output value, and wherein whether the processing rule is executed by the data processing application during processing of a particular data record depends directly or indirectly on a value in each of one or more data fields of the particular data record;

means for receiving execution information indicative of a number of times the processing rule was executed in connection with processing of the first set of data records;

means for obtaining a second set of data records;

means for analyzing values in one or more data fields of each of the data records in the second set, the analyzing including generating a profile of each of one or more of the data fields for the second set of data records, the profile of a data field characterizing the values in the data field;

means for obtaining a subsetting rule based on (i) the generated profiles and (ii) the execution information indicative of the number of times the processing rule was executed in connection with processing the first set of data records, the subsetting rule including an identification of a particular one of the data fields of the data records in the second set;

means for selecting a subset of data records from the second set of data records according to the subsetting rule, the selecting of the subset of data records being based on values in the particular one of the data fields; and

means for processing the selected subset of data records using the data processing application.

26. A computer-implemented method for selecting data records to cause execution of a processing rule during testing of a data processing application, the method including:
- obtaining a first set of data records;
 - processing in a data processing system the first set of data records using a data processing application that includes a processing rule that operates on at least one input value and generates at least one output value, and wherein whether the processing rule is executed by the data processing application during processing of a particular data record depends directly or indirectly on a value in each of one or more data fields of the particular data record;
 - receiving execution information indicative of a number of times the processing rule was executed in connection with processing of the first set of data records;
 - obtaining a second set of data records;
 - characterizing values in one or more data fields of each of the data records in the second set;
 - determining an identification criterion that identifies a particular one of the data fields of the data records in the second set, the identification criterion being based on (i) the characterization of the values in the one or more data fields of each of the data records in the second set and (ii) the execution information indicative of the number of times the processing rule was executed in connection with processing the first set of data records;
 - applying the identification criterion to select one or more data records from the second set of data records based on values in the particular one of the data fields; and
 - processing the selected one or more data records using the data processing application.
27. The method of claim 26, including identifying a particular one of the data fields of the data records in the second set based on a cardinality of each of one or more of the data fields of the data records in the second set.
28. The method of claim 27, wherein there are multiple distinct values in the particular one of the data fields across the data records in the second set, and wherein selecting one or more data records from the second set of data records includes selecting at least one data record having each of the distinct values in the particular one of the data fields.
29. The method of claim 26, wherein the particular one of the data fields is a first particular data field, and including identifying a second particular one of the data fields of the data records in the second set as a second particular data field.
30. The method of claim 29, wherein selecting one or more data records from the second set of data records includes selecting the data records based on a combination of values in the first particular

data field and values in the second particular data field.

31. The method of claim 30, wherein there are multiple distinct values in the first particular data field across the data records in the second set and multiple distinct values in the second particular data field across the data records in the second set, and wherein selecting one or more data records from the second set of data records includes selecting the data records based on combinations of the distinct values in the first particular data field and the distinct values in the second particular data field.

32. The method of claim 26, wherein characterizing values in one or more data fields includes identifying a first data field having a value that indicates a relationship among two or more data records in the second set of data records.

33. The method of claim 32, wherein selecting one or more data records from the second set of data records includes: selecting a first data record; identifying one or more second data records related to the first data record based on the value in the first data field of the first data record; and selecting the one or more second data records.

34. The method of claim 32, including determining an accumulated value for the first data record based on a value in a second data field of the first data record and a value in the second data field for each of the two or more data records that are related to the first data record; generating a new field for the first data record; and populating the new field with the accumulated value.

35. The method of claim 34, wherein determining the accumulated value includes determining the accumulated value based on a sum of the value in the second data field of the first data record and the value in the second data field for each of the two or more data records that are related to the first data record.

36. The method of claim 34, wherein selecting one or more data records from the second set of data records includes selecting the one or more data records based on the accumulated value in the new field for each of one or more data records in the second set of data records.

37. A system for selecting data records to cause execution of a processing rule during testing of a data processing application, the system including:

- at least one processor configured to:

- obtain a first set of data records;

- process in a data processing system the first set of data records using a data

- processing application that includes a processing rule that operates on at least one input value

and generates at least one output value, and wherein whether the processing rule is executed by the data processing application during processing of a particular data record depends directly or indirectly on a value in each of one or more data fields of the particular data record;

receive execution information indicative of a number of times the processing rule was executed in connection with processing of the first set of data records;

obtain a second set of data records;

characterize values in one or more data fields of each of the data records in the second set;

determine an identification criterion that identifies a particular one of the data fields of the data records in the second set, the identification criterion being based on (i) the characterization of the values in the one or more data fields of each of the data records in the second set and (ii) the execution information indicative of the number of times the processing rule was executed in connection with processing the first set of data records;

applying the identification criterion to select one or more data records from the second set of data records based on values in the particular one of the data fields; and

process the selected one or more data records using the data processing application.

38. The system of claim 37, wherein the at least one processor is configured to identify a particular one of the data fields of the data records in the second set based on a cardinality of each of one or more of the data fields of the data records in the second set.

39. The system of claim 38, wherein there are multiple distinct values in the particular one of the data fields across the data records in the second set, and wherein selecting one or more data records from the second set of data records includes selecting at least one data record having each of the distinct values in the particular one of the data fields.

40. The system of claim 37, wherein the particular one of the data fields is a first particular data field, and including identifying a second particular one of the data fields of the data records in the second set as a second particular data field.

41. The system of claim 40, wherein selecting one or more data records from the second set of data records includes selecting the data records based on a combination of values in the first particular data field and values in the second particular data field.

42. The system of claim 41, wherein there are multiple distinct values in the first particular data field across the data records in the second set and multiple distinct values in the second particular data

field across the data records in the second set, and wherein selecting one or more data records from the second set of data records includes selecting the data records based on combinations of the distinct values in the first particular data field and the distinct values in the second particular data field.

43. The system of claim 37, wherein characterizing values in one or more data fields includes identifying a first data field having a value that indicates a relationship among two or more data records in the second set of data records.

44. The system of claim 43, wherein selecting one or more data records from the second set of data records includes: selecting a first data record; identifying one or more second data records related to the first data record based on the value in the first data field of the first data record; and selecting the one or more second data records.

45. The system of claim 43, wherein the at least one processor is configured to determine an accumulated value for the first data record based on a value in a second data field of the first data record and a value in the second data field for each of the two or more data records that are related to the first data record; generating a new field for the first data record; populating the new field with the accumulated value.

46. The system of claim 45, wherein selecting one or more data records from the second set of data records includes selecting the one or more data records based on the accumulated value in the new field for each of one or more data records in the second set of data records.

47. A non-transitory computer readable medium storing instructions for causing a computer system to execute the method of any one of claims 26-36.

48. A system for selecting data records to cause execution of a processing rule during testing of a data processing application, the system including:

- means for obtaining a first set of data records;

- means for processing in a data processing system the first set of data records using a data processing application that includes a processing rule that operates on at least one input value and generates at least one output value, and wherein whether the processing rule is executed by the data processing application during processing of a particular data record depends directly or indirectly on a value in each of one or more data fields of the particular data record;

- means for receiving execution information indicative of a number of times the processing rule was executed in connection with processing of the first set of data records;

- means for obtaining a second set of data records;

means for characterizing values in one or more data fields of each of the data records in the second set;

means for determining an identification criterion that identifies a particular one of the data fields of the data records in the second set, the identification criterion being based on (i) the characterization of the values in the one or more data fields of each of the data records in the second set and (ii) the execution information indicative of the number of times the processing rule was executed in connection with processing the first set of data records;

means for applying the identification criterion to select one or more data records from the second set of data records based on values in the particular one of the data fields; and

means for processing the selected one or more data records using the data processing application.