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 as a target data field;

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 target data field; 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 one of the data fields as the target data field based on a cardinality of the identified one of the data fields.

3. The method of claim 2, wherein 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.

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 target data field based on the classifying.

5. The method of claim 4, wherein the target data field 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 target data field.

6. The method of claim 1, wherein the subsetting rule identifies a first data field as a first target data field and a second data field as a second target data field.

7. The method of claim 6, 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 target data field and a second set of distinct values for the second target data field.

8. 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.

9. The method of claim 8, 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.

10. The method of claim 8, 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.

11. 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.

12. The method of claim 11, 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.

13. The method of claim 1, wherein obtaining the subsetting rule includes receiving the subsetting rule.

14. The method of claim 1, including providing the selected subset of data records to the data processing application.

15. 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.

16. A non-transitory computer-readable medium storing instructions for causing a computing system to select data records to cause execution of a processing rule during testing of a data processing application, the instructions causing the computing system 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 as a target data field;

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 target data field; and

process the selected subset of data records using the data processing application.

17. The non-transitory computer-readable medium of claim 16, wherein obtaining the subsetting rule includes formulating the subsetting rule, including identifying the one of the data fields as the target data field based on a cardinality of the identified one of the data fields.

18. The non-transitory computer-readable medium of claim 17, wherein 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.

19. The non-transitory computer-readable medium of claim 16, 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 target data field based on the classifying.

20. The non-transitory computer-readable medium of claim 19, wherein the target data field 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 target data field.

21. The non-transitory computer-readable medium of claim 16, wherein the subsetting rule identifies a first data field as a first target data field and a second data field as a second target data field.

22. The non-transitory computer-readable medium of claim 21, 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 target data field and a second set of distinct values for the second target data field.

23. The non-transitory computer-readable medium of claim 16, 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.

24. The non-transitory computer-readable medium of claim 23, 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.

25. The non-transitory computer-readable medium of claim 16, 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.

26. The non-transitory computer-readable medium of claim 25, 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.

27. The non-transitory computer-readable medium of claim 16, including providing the selected subset of data records to the data processing application.

28. The non-transitory computer-readable medium of claim 16, 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.

29. 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 as a target data field;

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 target data field; and

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

30. The computing system of claim 29, wherein obtaining the subsetting rule includes formulating the subsetting rule, including identifying the one of the data fields as the target data field based on a cardinality of the identified one of the data fields.

31. The computing system of claim 30, wherein 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.

32. The computing system of claim 29, 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 target data field based on the classifying.

33. The computing system of claim 32, wherein the target data field 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 target data field.

34. The computing system of claim 29, wherein the subsetting rule identifies a first data field as a first target data field and a second data field as a second target data field.

35. The computing system of claim 34, 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 target data field and a second set of distinct values for the second target data field.

36. The computing system of claim 29, 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.

37. The computing system of claim 36, 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.

38. The computing system of claim 29, 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.

39. The computing system of claim 38, 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.

40. The computing system of claim 29, including providing the selected subset of data records to the data processing application.

41. The computing system of claim 29, 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.

42. 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 as a target data field;

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 target data field; and

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

43. 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;

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

identifying a particular one of the data fields of the data records in the second set as a target data field 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;

selecting one or more data records from the second set of data records based on values in the target data field; and

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

44. The method of claim 43, including identifying a particular one of the data fields of the data records in the second set as a target data field based on a cardinality of each of one or more of the data fields of the data records in the second set.

45. The method of claim 44, wherein there are multiple distinct values in the target 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 at least one data record having each of the distinct values in the target data field.

46. The method of claim 43, wherein the target data field is a first target data field, and including identifying a second particular one of the data fields of the data records in the second set as a second target data field.

47. The method of claim 46, 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 target data field and values in the second target data field.

48. The method of claim 47, wherein there are multiple distinct values in the first target data field across the data records in the second set and multiple distinct values in the second target 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 target data field and the distinct values in the second target data field.

49. The method of claim 43, 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.

50. The method of claim 49, 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.

51. The method of claim 49, 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.

52. The method of claim 51, 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.

53. The method of claim 51, 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.

54. 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 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;

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

identify a particular one of the data fields of the data records in the second set as a target data field 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;

select one or more data records from the second set of data records based on values in the target data field; and

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

55. The system of claim 54, 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 as a target data field based on a cardinality of each of one or more of the data fields of the data records in the second set.

56. The system of claim 55, wherein there are multiple distinct values in the target 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 at least one data record having each of the distinct values in the target data field.

57. The system of claim 54, wherein the target data field is a first target data field, and including identifying a second particular one of the data fields of the data records in the second set as a second target data field.

58. The system of claim 57, 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 target data field and values in the second target data field.

59. The system of claim 58, wherein there are multiple distinct values in the first target data field across the data records in the second set and multiple distinct values in the second target 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 target data field and the distinct values in the second target data field.

60. The system of claim 54, 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.

61. The system of claim 60, 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.

62. The system of claim 60, 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.

63. The system of claim 62, 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.

64. A non-transitory computer readable medium storing instructions for causing a computer system to select data records to cause execution of a processing rule during testing of a data processing application, the instructions causing the computer system 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;

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

identify a particular one of the data fields of the data records in the second set as a target data field 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;

select one or more data records from the second set of data records based on values in the target data field; and

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

65. The non-transitory computer readable medium of claim 64, the instructions causing the computer system to identify a particular one of the data fields of the data records in the second

set as a target data field based on a cardinality of each of one or more of the data fields of the data records in the second set.

66. The non-transitory computer readable medium of claim 65, wherein there are multiple distinct values in the target 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 at least one data record having each of the distinct values in the target data field.

67. The non-transitory computer readable medium of claim 64, wherein the target data field is a first target data field, and including identifying a second particular one of the data fields of the data records in the second set as a second target data field.

68. The non-transitory computer readable medium of claim 67, 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 target data field and values in the second target data field.

69. The non-transitory computer readable medium of claim 68, wherein there are multiple distinct values in the first target data field across the data records in the second set and multiple distinct values in the second target 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 target data field and the distinct values in the second target data field.

70. The non-transitory computer readable medium of claim 64, 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.

71. The non-transitory computer readable medium of claim 70, 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.

72. The non-transitory computer readable medium of claim 70, the instructions causing the computer system 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.

73. The non-transitory computer readable medium of claim 72, 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.

74. 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 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 characterizing values in one or more data fields of each of the data records in the second set;

means for identifying a particular one of the data fields of the data records in the second set as a target data field 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 selecting one or more data records from the second set of data records based on values in the target data field; and

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