



(51) International Patent Classification:

G06F 17/30 (2006.01)

(21) International Application Number:

PCT/US2012/065921

(22) International Filing Date:

19 November 2012 (19.11.2012)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

13/310,783 4 December 2011 (04.12.2011) US

(71) Applicant: BEYONDCORE, INC. [US/US]; 983 Shoal Dr., San Mateo, CA 94404 (US).

(72) Inventors: SENGUPTA, Arijit; C/o Beyondcore, Inc., 983 Shoal Dr., San Mateo, CA 94404 (US). STRONGER, Brad, A.; C/o Beyondcore, Inc., 983 Shoal Dr., San Mateo, CA 94404 (US). CHRONIS, Griffin; C/o Beyondcore, Inc., 983 Shoal Dr., San Mateo, CA 94404 (US).

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

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(74) Agents: FARN, Michael, W. et al; Fenwick &amp; West LLP, Silicon Valley Center, 801 California Street, Mountain View, CA 94041 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

(54) Title: ANALYZING DATA SETS WITH THE HELP OF INEXPERT HUMANS TO FIND PATTERNS




FIG. 1

(57) Abstract: A combined computer/human approach is used to detect actionable insights in large data sets. Automated computer analysis used to identify patterns (e.g., possibly meaningful patterns or subsets within the data). These are presented to humans for feedback, where the humans may have little to no training in the statistical methods used to detect actionable insights. Feedback from the humans is used to improve the pattern detection and facilitate the detection of actionable insights.

## ANALYZING DATA SETS WITH THE HELP OF INEXPERT HUMANS TO FIND PATTERNS

### BACKGROUND

The present invention relates to automated data analysis with the help of potentially untrained humans. In one aspect, it relates to leveraging structured "crowdsourcing" to enhance the analysis of data to find actionable insights and patterns. In another aspect, it relates to completely automating the analysis of data with humans only potentially assisting with the analysis if appropriate or necessary.

Traditional data analysis suffers from certain key limitations. Such analysis is used in a wide variety of domains including Six Sigma quality improvement, fraud analytics, supply chain analytics, customer behavior analytics, social media analytics, web interaction analytics, and many others. The objective of such analytics is to find actionable underlying patterns in a set of data.

Many types of analytics involve "hypothesis testing" to confirm whether a given hypothesis such as "people buy more pizza when it is raining" is true or not. The problem with such analytics is that human experts may easily not know of a key hypothesis and thus would not know to test for it. Analysts thus primarily find what they know to look for. In our quality improvement work with Fortune 100 firms and leading outsourcing providers, we have often found cases where clear opportunities to improve a process were missed because the analysts simply did not deduce the correct hypothesis.

For example, in a medical insurance policy data-entry process, there were several cases of operators marking applicants as the wrong gender. These errors would often go undetected and only get discovered during claims processing when the system would reject cases such as pregnancy related treatment for a policy that was supposed to be for a man. The underlying pattern turned out to be that when the policy application was in Spanish, certain operators selected "Male" when they saw the word *Mujer* which actually means female. In three years of trying to improve this process, the analysts had not thought to test for this hypothesis and had thus not found this improvement opportunity. Sometimes analysts simply do not have the time or resources to test for all possible hypotheses and thus they select a small subset of the potential hypotheses to test. Sometimes they may manually review a small subset of data to guess which hypotheses might be the best ones to test. Sometimes they interview process owners to try to select the best hypotheses to test. Because each of these cases is subject to human error and bias, an analyst may reject key hypotheses

even before testing it on the overall data. Thus, failure to detect or test for the right hypotheses is a key limitation of traditional analytics, and analysts who need not be domain experts are not very good at detecting such hypotheses.

Another limitation of traditional data analysis is the accuracy of the analysis models. Because the analysis attempts to correlate the data with one of the proposed models, it is critically important that the models accurately describe the data being analyzed. For example, one prospective model for sales of pizza might be as follows: Pizza sales are often correlated with the weather, with sporting events, or with pizza prices. However, consider a town in which the residents only buy pizza when it is both raining and there is a football game. In this situation, the model is unable to fit the data and the valuable pattern is not discovered. In one aspect of our invention, humans could recognize this pattern and provide the insight to the computer system.

A third limitation of traditional analysis is that the analysis is subject to human error. For example, many analysts conduct statistical trials using software such as SAS, STATA, or Minitab. If an analyst accidentally mistypes a number in a formula, the analysis could be completely incorrect and offer misleading conclusions. This problem is so prevalent that one leading analysis firm requires all statistical analyses to be performed by two independent analysts and the conclusions compared to detect errors. Of course, this is just one way in which humans can introduce error into the broad process of bringing data from collection to conclusion.

Finally, because humans cannot easily deal with large volumes of data or complex data, analysts often ignore variables they deem less important. Analysts may easily accidentally ignore a variable that turns out to be key. During an analysis of a credit card application process, it was found that the auditors had ignored the "Time at current address" field in their analysis as it was thought to be a relatively unimportant field. However, it turned out that this field had an exceptionally high error rate (perhaps precisely because operators also figured that the field was unimportant and thus did not pay attention to processing it correctly). Once the high error rate was factored in, this initially ignored field turned out to be a key factor in the overall analysis. Analysts also sometimes initially explore data to get a "sense of it" to help them form their hypotheses. Typically, for large datasets, analysts can only explore subsets of the overall data to detect patterns that would lead them to the right hypotheses or models. If they accidentally look at the wrong subset or fail to review a subset with the clearest patterns, they may easily miss key factors that would affect the accuracy of their analysis. Analysts may similarly exclude sets of data from the analysis

because they believe these are outliers. However, sometimes truly useful data may be incorrectly excluded by human analysts.

On the other hand, an emerging best practice in the world of business analytics is the practice of "crowdsourcing." This refers to tapping a large set of people (the "crowd") to provide insight to help solve business issues. For example, a customer might fill out a comment card indicating that a certain dress was not purchased because the customer could not find matching shoes. This can be a very valuable insight, but the traditional collection procedure suffers from several problems.

The first step in crowdsourcing is undirected social idea generation. Employees, customers, and others submit ideas and patterns that they have identified. Of course, any pattern that is not noticed by a human is not submitted and is therefore not considered in the analysis.

The next step is for someone to sort and filter all the submitted ideas. Because there are a large volume of suggestions, and it is impossible to know if the suggestions are valuable without further research, someone must make the decision on which ideas to follow up on. This can be based on how many times an idea is submitted, how much it appeals to the people sorting the suggestions, or any number of methods. The issue is that good ideas may be rejected and never investigated.

Once the selected ideas are passed to an analyst, he or she must decide how to evaluate the ideas. Research must be conducted and data collected. Sometimes the data is easily available, for example, if a customer suggests that iced tea sells better on hot days, the sales records can be correlated with weather reports. Sometimes the data must be gathered, for example, if a salesman thinks that a dress is not selling well due to a lack of matching shoes, a study can be performed where the dress is displayed with and without clearly matching shoes and the sales volumes compared. However, sometimes it is impossible to validate a theory because the corresponding data is not available.

Finally, the analysis is only as good as the analyst who performs it in the first place. An inexperienced analyst often produces much less useful results than an experienced analyst even when both work on the same data.

Thus there is a need for a solution which takes the strengths of the computer and the strengths of the humans and leverages both in a scalable manner. Such a solution could increase the effectiveness of analytics by decreasing the impact of human errors and human inability to select the correct hypotheses and models.

Further, there is a need for a scalable approach to crowdsourcing which does not suffer from the limitations of traditional crowdsourcing described above.

On the other hand, automated analysis also suffers from certain limitations. The software may not see that two different patterns detected by it are actually associated or be able to detect the underlying reason for the pattern. For example, in the policy data entry example described above, an automated analysis could detect that Spanish forms had higher error rates in the gender field but automated analysis may not be able to spot the true underlying reason. A human being however may suggest checking the errors against whether or not the corresponding operator knew Spanish. This would allow the analysis to statistically confirm that operators who do not know Spanish exhibit a disproportionately high error rate while selecting the gender for female customers (due to the *Mujer* = male confusion).

## BRIEF DESCRIPTION OF THE DRAWINGS

The preferred embodiments of the invention will hereinafter be described in conjunction with the appended drawings provided to illustrate and not to limit the invention, wherein like designations denote like elements, and in which:

Fig. 1 is block diagram illustrating a combined computer/human approach to finding patterns and other actionable insights in large data sets.

## DESCRIPTION OF PREFERRED EMBODIMENTS

### Introduction

Fig. 1 is block diagram illustrating a combined computer/human approach to finding patterns and other actionable insights in large data sets. For simplicity, most of the following discussion is in the context of finding meaningful patterns, but the principles illustrated can also be applied to identify other types of actionable insights. Steps 110 and 120 are largely based on automatic computer analysis. Step 130 is largely based on human analysis (e.g., crowdsourcing), preferably by relatively untrained humans. Optional step 140 is a more formal analysis, preferably by trained analysts, to test and refine the findings and suggestions of steps 110-130.

In this context, "untrained" means little to no training in the statistical principles underlying the search for actionable insights. The term "statistically untrained" may sometimes be used. Thus, in a conventional approach, statistical analysts review the data, form hypotheses and design and run statistically significant experiments to test their

hypotheses. These statistical analysts will generally be "trained humans" or "statistically trained humans." On the other hand, consider a case where the process being examined is a loan underwriting process. Feedback may be solicited from humans ranging from data entry operators to those making final decisions on loan applications. These will generally be "untrained humans" or "statistically untrained humans" because they are not providing feedback on the statistical aspect of the search for actionable insights. Note that the term "untrained human" does not mean that these humans are unskilled. They may be highly trained in other areas, such as loan evaluation. They may even have enough training in statistics to play the role of a statistical analysis; they just are not playing that role in this context.

Steps 110 and 120 are the automatic analysis of large data sets and the automatic detection of potentially valuable and meaningful patterns within those data sets. We have previously disclosed multiple approaches to automatically analyzing data to detect underlying patterns and insights. Examples include U.S. Patent 7,849,062 "Identifying and Using Critical Fields in Quality Management" that disclosed means to automatically detect underlying error patterns in data processing operations as well as pending patent application PCT/US2011/033489 "Identifying and Using Critical Fields in Quality Management" that disclose additional approaches to automatically analyzing data to detect underlying patterns. While some of these inventions were described in the context of data processing or human error patterns detection, the underlying methods are also applicable to a broad range of analytics. In U.S. Patent Application No. 13/249,168 "Analyzing Large Data Sets to Find Operator Deviation Patterns," we specifically disclosed approaches that allowed the automatic detection of subsets of data with high p-values indicating the high likelihood that the specific subset contained some underlying patterns and that the corresponding data distribution was unlikely to have been random. Thus, the underlying patterns have a higher chance of leading to meaningful actionable insights. Another approach to detecting actionable insights involves analyzing the predictive determinism of a set of variables or variable combinations, as measured by the coefficient of determination  $R^2$ , or goodness of fit against a regression model such as a piece-wise linear or "kitchen-sink" regression. These approaches can be applied to analyses including but not limited to customer segmentation (psychographics), sales analysis, marketing campaign optimization, demand forecasting, inventory / resource / supply chain optimization, assortment / product mix optimization, causal analysis, fraud detection, overtoiling detection, and risk analysis. All of the foregoing are incorporated by reference herein.

The output of such automated analysis 110/120 can be further enhanced by the addition of manual feedback 130. Such feedback can be provided by statistically trained humans, however, certain types of extremely valuable feedback can be provided by statistically untrained humans. For example, a company's employees, customers, suppliers or even interested humans without special knowledge / experience may be able to provide valuable feedback that can enhance the automated analysis 110/120.

For example, in the policy data entry example described above, an automated analysis 110/120 could detect that Spanish forms had higher error rates in the gender field but automated analysis may not be able to spot the true underlying reason. A human being however may suggest 130 checking the errors against whether or not the corresponding operator knew Spanish. As indicated by the feedback arrow 135, this would allow the analysis 110/120 to statistically confirm that operators who do not know Spanish exhibit a disproportionately high error rate while selecting the gender for female customers (due to the *Mujer* = male confusion). In this way, actionable insights can be iteratively developed through a combination of computer analysis and statistically untrained human feedback.

One goal here is to minimize the need for expert knowledge, such as deep understanding of statistics, so that the scope of potential crowdsourcing contributors 130 is as broad as possible. At the same time, an additional goal is to make the opportunities for crowdsourcing feedback 135 sufficiently structured in nature, such that the overall process can be as automated as possible and does not require subjective human evaluation or processing of the crowdsourced feedback. A final optional goal is tying the crowdsourced feedback and the automated analytics tightly and interactively to the available data so that the analysis produces actionable insights backed by statistically valid evidence.

#### Automated analysis

Various types of automated analysis have been described previously by the inventors. For example, in the context of document processing by operators, one goal may be to find documents that are similar in some way in order to identify underlying patterns of operator behavior. A search can be conducted for segments of the data which share as few as one or more similar field or parameter values. For example, a database of loan applications can be searched for applicants between 37 and 39 years of age. Any pair of applications from this sample might be no more similar than a randomly chosen pair from the population. However, this set of applications can be statistically analyzed to determine whether certain loan officers are more likely to approve loans from this section of the population.

Alternatively, it may not be necessary to find even one very similar parameter. Large segments of the population may be aggregated for analysis using criteria such as "applicants under 32 years old" or "applicants earning more than \$30,000 per year." Extending this methodology one step further, a single analysis can be conducted on the sample consisting of the entire population.

In addition, it is possible to analyze sets of data which do not contain all of the information that the operators use to make decisions. In the case of loan applications requiring a personal interview, it would be very hard to conduct a controlled experiment that includes the personal interview. It would also be difficult to search for "similar" interviews. However, we can still search for applications with some parameters similar, and aggregate the statistics across all interviews. It may not be possible to identify any single loan decision as incorrect or suspect, but if, for example, among applicants aged 26-28, earning over \$32,000, one loan officer approves 12% of loans and another approves 74% of loans, there may be training or other issues.

These methods can be combined to find a diverse variety of samples to analyze. A sample might consist of the documents with each field similar to a given value for that field, or it might comprise the set of all the documents. In addition, some fields may be restricted to a small or large range, where other fields have no restriction. Each sample may be analyzed with statistical methods to determine whether operators are processing documents consistently.

There are several statistical hypothesis tests which may be appropriate for making this determination. If the output of the process is binary, such as a loan approval, and the number of documents in the sample under analysis is small, a test such as Fisher's Exact Test may be used. If the output is a number, such as a loan interest rate, and the sample is large, a Chi-Square Test may be used. These tests can be used to determine whether one operator is producing significantly differing output from the remainder of the operators. Alternately, the operators can be split into two groups and these tests can be used to determine whether the operators in the two groups are producing significantly differing output. All possible splits can be analyzed to find the one with the highest statistical significance. Alternately, these tests can be used to determine simply whether the distribution of operator output for this sample is significantly more unusual than what would be expected under the null hypothesis, i.e., all operators making decisions in the same manner.

If numerous statistical tests are conducted, it is expected that some of them will be statistically significant, even without any underlying pattern. It is important to search for p-



values which are more extreme than would normally be sought. For example, if 1000 tests are conducted, we could require a p-value of 0.00005 rather than the typical 0.05.

Alternately, we can split the data into two sets of data: a training set and a testing set. We can conduct a large number of tests on the training data, but may find that our lowest p-values are not statistically significant due to the large number of tests conducted. We can then use the results to construct new hypotheses and design a small number of new tests. These new tests can be conducted on the testing data set. Because only a few tests are being conducted, we would not need very extreme p-values to achieve significance. Alternately, we can use the results as a starting point for any other review process, including supervisor review of indicated historic documents. Rules can potentially also be created to automatically flag documents from this segment of the population, as they are processed, for additional review.

Another method for computing the statistical significance of complicated test statistics is as follows. We are testing against the null hypothesis that all operators behave in the same manner. Disproving this null hypothesis means there is some statistically significant underlying pattern to the behavior of the operators. For statistics where operators are separated into multiple groups under a grouping plan, we can randomly assign operators into groups repeatedly under multiple different grouping plans and re-compute the test statistic for each grouping plan. If the value for a specific grouping plan is higher than the value for 95% of randomized grouping plans then we have 95% confidence that our null hypothesis was incorrect. Of course, we cannot simply compute many random grouping plans and assert that the top few grouping plans are statistically significant. However, we can identify a possibly significant grouping plan by doing this for the training dataset, and see if that grouping plan is again in the best 5% of random grouping plans for the testing data set.

It should be noted that a statistical hypothesis test can be very useful for showing that one or more operators produce different output (or a different output distribution) for documents from the same section of the population. However, it may be more valuable to find sections of the population where the operator output difference is large, rather than merely statistically significant. Metrics other than statistical p-value can be used to determine which population sections require further study. One such metric is related to the variance in the means of the operators output. Because we only have access to a sample of the data, we typically cannot compute the actual means. We can instead compute an estimate of each of the means and use this to calculate an estimate of the variance in the means.

In a stable process where there were no deviations from the norm, the variance would be significantly lower than in a process with patterns of deviations from the norm. Any of

these metrics, or others, can be used as the basis of a hill climb or other local search method to identify interesting samples of the population that would be most useful to analyze to detect underlying patterns of deviations from norms or fragmented norms. A key property of these metrics is that they are highest for the section of the document population that actually represents the variance in operator behavior. For example, if one operator is not approving loans from males aged 20-30, the metric should be higher for "males aged 20-30" than for "males aged 20-50" and "people aged 20-30."

Local search methods operate by considering a given sample, and repeatedly modifying it with the goal of raising the metric. This continues until the metric is higher for the sample under consideration than for any nearby samples (a local optimum). The notion of proximity is complex for samples of the sort we are discussing. The "modify" step in the algorithm will change the restrictions defining the current sample. This can consist of widening or tightening the restriction on one field, or adding a restriction on a new field, or removing the restriction on a restricted field. For example, if we consider a sample consisting of "Loan applications from females aged 30-40" and calculate the metric to be X, we could then calculate the metric for "females", "females aged 30-50", "females aged 20-40", "people aged 30-40", and others. Each of these metrics will be compared to X and the search algorithm will continue. Such a search can also be used to detect non-obvious relationships between different variables. For example, females may in general have a slightly higher interest rate than usual, and people aged 18 to 25 may have a higher interest rate than usual, but females aged 18 to 25 may actually have a lower interest rate than usual. Such a pattern would have been ignored if we just used a regression model that did not specifically include a term for the combination of gender and age.

Because the metrics are highest for samples with acute variances, samples obtained using parameter values which are responsible for the unusual behavior will have the highest scores. Much larger and much smaller samples will have lower scores. As the search algorithm runs, the sample under consideration will "evolve" to contain the features that are causing the discrepancy in operator processing while not containing unrelated random information. Of course, the search will cease on one local maximum. If the local search is repeated multiple times from random starting samples, many samples with peak metrics can be identified in the data. The same kind of analysis, including local search, can be used to identify the variables or combination of variables that have the highest predictive determinism for specified outcome variable. The outcome variable can be specified by an expert or the software can try different variables as potential outcome variables and then

select the outcome variable based on specified selection criteria. For example, the variable that exhibits the highest signal to noise may be selected as the outcome variable. The outcome variable may also be calculated from a set of variables.

The examples above were given in the context of forming hypotheses for patterns of operator behavior, but they can also be used to form hypotheses for other types of analysis. For example, one type of analysis may involve the creation of reports such as charts or graphs based on the variables or combinations of variables with the highest predictive determinism. These hypotheses can then be further qualified 130 by humans.

#### Human social ideation

Referring to Fig. 1, human feedback 130 is used to improve the hypotheses identified by the automated analysis 110/120. Multiple forms of directed crowdsourced or social feedback can be supported. Examples include the following.

Voting of auto-detected patterns: Humans may simply review the auto-detected patterns or subsets of data with high p-values and vote that the specific pattern or subset is worth further exploration. The higher the number of votes a pattern gets, the more actionable or worthy of further exploration the pattern might be. Similarly they may vote on the best outcome variables for predictive determinism analysis. They may also vote on variables or variable combinations that are best to generate analysis reports by.

Tagging of auto-detected patterns: Humans may also tag the patterns or subsets with comments. For example, in an invoice processing scenario, certain operators might incorrectly process debits as credits. This error would show up in different ways. First, the amount for the line item would be positive instead of negative. Second, the transaction type would be marked incorrectly. And finally, the total amount for the invoice would be incorrect. While automated analysis might detect that the three patterns are highly correlated it might not have sufficient information to reveal that there is a causal relationship between the patterns. One or more humans however may tag the three different error patterns as part of a broader "debit/credit confusion" pattern. This would help the automated analysis detect the fact that a single underlying problem, operators confusing debits and credits, is the root cause behind these multiple patterns. Another tagging example could occur for an automated analysis that revealed that a certain bank was issuing very few loans below \$10,000 and that this pattern had significant statistical evidence of being significant. A human might however know that the specific bank only serves multi-millionaires and thus rarely received loan applications for small amounts. The human could thus tag this pattern as not worth exploring

due to this reason. If sufficient humans tagged the pattern the same way, the automated analysis may reduce the importance of the pattern despite the high statistical evidence.

Propose hypotheses : The analytics may reveal patterns but due to the lack of understanding of the complex real world systems, algorithms may not detect the right corresponding hypotheses. For example, the analysis may reveal that something statistically significant is happening which is causing a significantly lower sale of certain dresses in certain shops as opposed to other shops even though the dresses were displayed the same way in all stores on identical mannequins. A customer may point out that the dress material displays certain attractive characteristics when seen under florescent light and not under incandescent light. This would be an example of a hypothesis that an automated analysis probably would not identify and even human experts may have easily missed. However, given a specific pattern to focus on as a starting point, at least one of a sufficiently large number of crowdsourced helpers may detect this key pattern.

Filter / search data to find new slices with high p-values : Automated analysis might leverage various heuristics such as "hill climb" to detect the subsets with the highest p-values. However, humans, especially customers and employees, because of their unique understanding of the broader context may be able to find subsets of data with high p-values that automated analysis did not detect. Humans may also realize that certain subsets were actually related and propose more complex subsets that would have even higher p-values. Additionally, because of heuristics like bucketing, the automated analysis may have somewhat imprecisely defined the subset and unnecessarily included/excluded data points in the subset that did not/did relate to the underlying pattern in question. Humans may define the subset more precisely, either including related data points or excluding unrelated data points to increase the p-values. For example, the system might detect an unusual volume of sales between \$20 and \$30 during the March 1-15 time period. A customer might remember a promotion of a free gift with purchases over \$25 during Feb 25 to Mar 12 and suggest this as a new subset to analyze, leading to an even higher p-value.

Propose external variables or datum to consider : A key limitation of automated analysis is the lack of awareness of the physical world or overall context. Humans may easily recommend the inclusion of additional variables, the inclusion of which simplifies or enables the detection of patterns. For example, if the automated analysis was evaluating the sale of pizzas, humans might suggest the inclusion of key causal variables such as the dates on which football games are held, or the local rainfall rates as these variables significantly affect the sale of home-delivered pizza. Similarly humans may simply provide additional

specific information such as "This specific shop uses incandescent lights" rather than suggest an external variable to consider.

Suggest fields to combine during analysis: Certain patterns may be relatively complex, such as "if variable A is equal to x and variable B is greater than y but variable C is not equal to z, then a specific pattern is observed." Such complex patterns may be difficult for automated analysis to detect short of expensive brute force analysis of an enormous number of possible scenarios. Humans, because of their enhanced understanding of the context, can more easily suggest such patterns.

Suggest breaking existing data into finer grained fields: Certain fields may represent overly aggregated data which hides underlying patterns. For example, if sales data is aggregated by day, a user may suggest that sales in the morning and in the evening should be tracked separately because different types of customers visit the shop during the morning as opposed to the evening and they exhibit different sales behavior patterns.

Suggest type of regression: Humans may have an instinct for the shape of the hidden data distribution. For example, humans may be asked to vote on whether the underlying pattern is linear, exponential, etc. They may also suggest combining certain variables during the analysis as specified in f above. In each of these cases, they are essentially suggesting the type of regression that the automated analysis should use.

Suggest experiments to detect or confirm patterns: In some cases, the humans may be aware of a pattern that cannot be confirmed from just the available data. For example, if a dress was not selling because customers could not imagine what kind of shoe they could wear with it, merely analyzing existing data may not be sufficient. However, human feedback may suggest that this hypothesis be tested by setting up floor displays with the specific dress and corresponding shoes or selling the dress and matching shoes together as a package. The results of this experiment would offer data that could confirm this hypothesis.

The previous section talks about auto-detected patterns or auto-detected subsets of data with high p-values. However, this method may be applied to other forms of automated, assisted, or manual data analysis as well. For example, there is no reason to believe that such social feedback would not be useful to an expert analyst performing a completely manual data analysis.

#### Collection of human feedback

Although feedback can be solicited as free-form text, there are several ways that we can structure the collection of feedback from customers and others. Structured as opposed to

free-form feedback allows easier automated understanding of the feedback as well as enhanced clustering of feedback to determine cases where multiple humans have essentially provided the same feedback.

One method for collecting structured feedback involves having users select each word in a sentence from a drop-down of possible words. In this way they can construct a suggestion, comment, or other insight such as "I would purchase more shoes if they were red." Each of the nouns and verbs can be altered but the sentence structure remains easy to analyze. The user could choose from insight templates such as "I would X if Y," "I feel X when Y," "I enjoy X when Y," etc.

For cases where the feedback involves filtering / searching data to find new slices with high p-values, the structured interface can be similar to standard advanced search functionality. The criteria specified by the human can be immediately tested on all the data or a selected subset of the data and the p-value measured.

For cases where the feedback involves suggesting outcome variables for predictive determinism analysis, or actually specifying the best variables for analyzing and reporting on the data, the structured interface can simply be a list of variables that the human can select from. In complex situations, the structured interface can allow the specification of mathematical formulas that combine selected variables into a complex outcome or analysis/reporting variable. The variables specified by the human can be immediately tested on all the data or a selected subset of the data and the predictive determinism or goodness of fit can be measured.

Another way to accept structured feedback is to ask the users to construct their sentence using a restricted language of selected nouns, verbs, and adjectives. These can be automatically analyzed by software algorithms such as statistical aggregation, Markov chains, and others to detect patterns.

If no other option allowed the user to express herself fully, she could compose her thoughts in free-form text. However, instead of having this text interpreted by humans, it could be analyzed by computer algorithms such as statistical aggregation, Markov chains, and others as described above.

Humans may be provided financial or other rewards based on whether their feedback was useful and unique. For example, in the filtering case, a user might be rewarded based on the feedback's usefulness, namely how much better the p-value of their specified subset was than the average p-values of the top 10 subsets previously detected by the software automatically or with the help of humans. A uniqueness criterion may also be easily applied

to the reward formula such that a higher reward would be paid if the human-specified subset differed significantly from previously identified subsets. The uniqueness of a user specified set  $N$  as compared to each of the previously identified sets  $S_i$  may be determined by a formula such as the following:  $(\text{Number of elements in } N - \text{Number of element in } N \text{ intersect } S_i) / (\text{Number of element in } N \text{ intersect } S_i)$ . Other uniqueness and usefulness criteria might be applied instead or in addition.

For feedback involving regression models or combinations of fields to be used in the model, a very similar approach combining usefulness and uniqueness can be used. Usefulness can be determined by the improvement in the "fit" of the model or increase in the predictive determinism of a combination of variables, while uniqueness can be determined by whether a substantially similar model has already been submitted previously or detected automatically.

Alternate approaches to rewards may include the following for cases where humans are tagging or voting for a pattern. The first person to tag a pattern with a given phrase might be rewarded based on how many other users also tagged the same pattern with the same phrase. This motivates users to tag with the phrases that they think other users will tag with. Even a software algorithm that attempted to "game" this system would, if successful, provide valuable insight. Given that users would not know what phrases a given pattern has already been tagged with, or even whether a pattern has already been tagged, it would be difficult for a user to predictably game such a system to get unwarranted rewards. Rewards can be restricted to tags that are uniquely popular for this pattern, to avoid the possibility every pattern getting tagged with a trivial tag. Alternately, the reward can be reduced if a user provides lot of tags. Thus, users would have an incentive to provide a few tags that are good matches for the data rather than a lot of less useful tags in the hope that at least one of the tags would be a good match.

Most reward-incented systems rely on rewards which are delayed in time with respect to the feedback offered by users. Because this system as described can measure p-values interactively, rewards can be immediately awarded, significantly improving the perceived value of participating in the system and increasing participation.

The structured human feedback process may be transformed into games of various sorts. Various games related to human-based computation have been used to solve problems such as tagging images or discovering the three dimensional shape of protein structures. This is just one example of how using automated analysis to create a good starting point and then allowing a framework where different humans can handle the tasks most suited to their

interests and abilities, can be more effective than either just automated or just expert manual analysis.

Existing approaches can be further improved in a number of ways. For example, one embodiment taps a human's social knowledge, something much harder for computers to emulate than specific spatial reasoning. Moreover, we tap the social knowledge in a structured machine-interpretable manner which makes the solution scalable. Humans excel at graph search problems such as geometric folding (or chess-playing) where there are many options at each step. Today, this gives people an advantage in a head-to-head competition, but with rapid advances in technology and falling costs, computers are rapidly catching up. In fact, computer algorithms are now widely considered to outperform humans at the game of chess. However, no amount of increased processor speed will enable a computer to compete in the arena of social cognizance and emotional intelligence. Socialization comes naturally to humans and can be effectively harnessed using our methods.

Additionally, various embodiments can be non-trivially reward based. By tying a tangible payment to the actual business value created, the system is no longer academic, but can encourage users to spend significant amounts of time generating value. Additionally, a user who seeks to "game" the system by writing computer algorithms to participate is actually contributing to the community in a valid and valuable way. Such behavior is encouraged. This value sharing approach brings the state of the art in crowdsourcing out of the arena of research papers and into the world of business.

Finally, some approaches allow humans to impact large aspects of the analysis, not just a small tactical component. For example, when a human suggests the inclusion of an external variable or identifies a subset with high p-value, they can change the direction of the analysis. Humans can even propose hypotheses that turn out to be the key actionable insight. Thus, unlike in the image tagging cases, humans are not just cogs in a computer driven process. Here, humans and computers are synergistic entities. Moreover, even without explicit collaboration, each insight from a human feeds back into the analysis and becomes available to other humans to build on. For example, Andy may suggest the inclusion of an external variable which leads Brad to detect a new subset with extremely high p-value, which leads Darrell to propose a hypothesis and Jesse to propose a specific regression model which allows the software to complete the analysis without expert human intervention. Thus, the human feedback builds exponentially on top of other human feedback without explicit collaboration between the humans.



Some humans may try to submit large volumes of suggestions hoping that at least one of them works. Others may even write computer code to generate many suggestions. As long as the computation resources needed to evaluate such suggestions is minimal, this is not a significant problem and may even contribute to the overall objective of useful analysis. To reduce the computational cost of the evaluation of suggestions, such suggestions may first be tested against a subset of the overall data. Suggestions would only be incorporated while analyzing the overall data if the suggestion enabled a significant improvement when used to analyze the subset data. To further save computation expenses, multiple suggestions evaluated on the subset data may be combined before the corresponding updated analysis is run on the complete data. Additionally, computation resources could be allocated to different users via a quota system, and users could optionally "purchase" more using their rewards from previous suggestions.

#### Feedback loop

Once the feedback is received 135, the initial automated analysis 110/120 may be re-run. For example, if the humans suggested additional external data, new hypotheses, new patterns, new subsets of data with higher p-values, etc., each of these may enable improved automated analysis. After the automated analysis is completed in light of the human-feedback, the system may go through an additional human-feedback step. The automated-analysis through human feedback cycle may be carried out as many times as necessary to get optimal analysis results. The feedback cycle may be terminated after a set number of times or if the results do not improve significantly after a feedback cycle or if no significant new feedback is received during a given human feedback step. The feedback cycle need not be a monolithic process. For example, if a human feedback only affects part of the overall analysis, that part may be reanalyzed automatically based on the feedback without affecting the rest of the analysis.

As the analysis is improved based on human feedback, a learning algorithm can evaluate which human feedback had the most impact on the results and which feedback had minor or even negative impact on the results. As this method clearly links specific human feedback to specific impacts on the results of the analysis, the learning algorithms have a rich source of data to train on. Eventually, these learning algorithms would themselves be able to suggest improvement opportunities which could be directly leveraged in the automated analysis phase. The automated analysis methods could themselves be improved based on learning algorithms analyzing which automatically detected patterns received the most

positive human feedback and which automatically detected patterns were rejected by the humans.

The human feedback patterns could also be analyzed to detect deterministic patterns that may or may not be context specific. For example, if local rainfall patterns turn out to be a common external variable for retail analyses, the software may automatically start including this data in similar analyses. Similarly, if humans frequently combine behavior patterns noticed on Saturdays and Sundays to create a higher p-value pattern for weekends, the software could learn to treat weekends and weekdays differently in its analyses.

The software may also detect tags that are highly correlated with (usually paired with) each other. If a pattern is associated with one of the paired tags but not the other, this may imply that the humans simply neglected to associate the pattern with the other tag, or it may be a special rare case where the pattern is only associated with one of the usually paired tags. The software can then analyze the data to detect which of the two cases has occurred and adjust the analysis accordingly.

This overall feedback loop may occur one or more times and may even be continuous in nature where the analysis keeps occurring in real time and users simply keep adding more feedback and the system keeps adjusting accordingly. An example of this may be a system that predicts the movement of the stock market on an ongoing basis with the help of live human feedback.

During the crowdsourcing phase, certain data will be revealed to the feedback crowd members. Companies may be willing to reveal different amounts and types of data to employees as opposed to suppliers or customers or the public at large. Security/privacy can be maintained using different approaches, including those described in U.S. Patent No. 7,940,929 "Method For Processing Documents Containing Restricted Information" and U.S. Patent Appls. 13/103,883 "Shuffling Documents Containing Restricted Information" and 13/190,358 "Secure Handling of Documents with Fields that Possibly Contain Restricted Information". All of the foregoing are incorporated by reference herein.

#### Further Analysis

Once the automated analysis with human feedback is completed, the data could be presented to expert analysts 140 for further enhancement. Such analysts would have the benefit of the following:

- lists of hypotheses detected automatically as well as proposed by humans;

- results of how well the data fit various regression models detected automatically as well as proposed by humans;
- specific subsets of data with high p-values, corresponding to automatically or manually detected patterns, and corresponding manually proposed causal links;
- votes and tags indicating agreement from communities such as customers or employees; and
- other valuable context information

Such information significantly ameliorates some of the key limitations of manual expert analysis such as picking the wrong hypotheses, the wrong models, ignoring key variables, reviewing the wrong subsets, etc.

The analyst's responsibilities can also be restricted to tasks such as slightly changing models, etc. or improving the way the data is analyzed rather than having to write complex code from scratch or figuring out which data sources need to be included in the analysis. By reducing the complexity and the "degrees of freedom" of the work the analyst has to perform, we significantly reduce the risk of human error or the impact of an analyst's experience on the final results. This may also enable superior analysis with lower cost analysts.

Given the nature of the automated analysis, the structured nature of the crowdsourced feedback, and the minimal optional involvement of expert analysts, such an analysis can be carried out much faster, at lower overall cost and higher overall accuracy and effectiveness than traditional methods.

Given the report-writing flexibility and freedom that analysts enjoy under traditional methods, it can be difficult to create scalable user-friendly reports with drill-down, expand-out, context-aware features and context specific data details. In essence, when an analyst writes custom code or analysis formulae to create analyses, the reports themselves have to be custom in nature and are difficult to build automatically without manual customization. However, the methodology specified above can restrict the expert analyst to configure, not customize. Due to the nature of the automated analysis, the structured feedback, and the limited expert configuration, the software solution is fully aware of all aspects of the report context and can automatically generate a rich context specific report with drill-down, expand-out, context specific data capabilities. The software can also combine measures of statistical soundness, predictive determinism, and/or P-values into a simple measure of statistical importance. This may simply be a weighted average of these and other statistical measures or other similar means of combining different metrics into a single composite one. The measure

of statistical importance may be presented relative to the variable combination or pattern with the highest overall statistical importance and may be expressed as a percentage.

One example of a rich context specific report that the software may automatically generate includes an animation that systematically walks a human through the key insights in the data without any human intervention. The methods described in this document can be used to determine the most useful subset of the data to analyze, the set of variables and or hypotheses and or regressions to analyze by, and the set of variables to report by. The animation can walk the human through different graphs based on different variables and variable combinations that have high predictive determinism and/or high P-values. The animation can also highlight different actionable insights such as women aged 18 to 25 behave significantly differently than women in general or 18 to 25 year olds in general. The animation can also exclude or highlight cases where a pattern is not statistically significant, has low predictive determinism or a low P-value and thus avoid humans wasting time on patterns that are not statistically sound. Note that while what data and graphs are presented is being determined based on complex and powerful statistics, these details can be completely hidden from the human who may just be presented a set of simple graphs and insights. In this manner, a rich statistically sound report can be presented solely through examples, and the user would not need to know complex statistics to benefit from the report.

The system, as described in the present invention or any of its components, may be embodied in the form of a computer system. Typical examples of a computer system include a general-purpose computer, a programmed microprocessor, a micro-controller, a peripheral integrated circuit element, and other devices or arrangements of devices that are capable of implementing the steps that constitute the method of the present invention.

The computer system comprises a computer, an input device, a display unit and the Internet. The computer comprises a microprocessor. The microprocessor can be one or more general- or special-purpose processors such as a Pentium®, Centrino®, Power PC®, and a digital signal processor. The microprocessor is connected to a communication bus. The computer also includes a memory, which may include Random Access Memory (RAM) and Read Only Memory (ROM). The computer system also comprises a storage device, which can be a hard disk drive or a removable storage device such as a floppy disk drive, optical disk drive, and so forth. The storage device can also be other similar means for loading computer programs or other instructions into the computer system. The computer system also includes one or more user input devices such as a mouse and a keyboard, and one or more output devices such as a display unit and speakers.

The computer system includes an operating system (OS), such as Windows, Windows CE, Mac, Linux, Unix, a cellular phone OS, or a proprietary OS.

The computer system executes a set of instructions that are stored in one or more storage elements, to process input data. The storage elements may also hold data or other information as desired. A storage element may be an information source or physical memory element present in the processing machine.

The set of instructions may include various commands that instruct the processing machine to perform specific tasks such as the steps that constitute the method of the present invention. The set of instructions may be in the form of a software program. The software may be in various forms, such as system software or application software. Further, the software may be in the form of a collection of separate programs, a program module with a larger program, or a portion of a program module. The software might also include modular programming in the form of object-oriented programming and may use any suitable language such as C, C++ and Java. The processing of input data by the processing machine may be in response to user commands to results of previous processing, or in response to a request made by another processing machine.

While the preferred embodiments of the invention have been illustrated and described, it will be clear that it is not limited to these embodiments only. Numerous modifications, changes, variations, substitutions and equivalents will be apparent to those skilled in the art, without departing from the spirit and scope of the invention, as described in the claims.

## WHAT IS CLAIMED IS:

1. A method implemented by a computer system for automatically identifying actionable insights in a data set, the method comprising the computer system performing the steps of:  
analyzing data within the data set, wherein such analysis occurs without human intervention and without use of predefined rules;  
identifying potentially valuable patterns based on the data analysis; and  
developing actionable insights based on the identified potentially valuable patterns.
2. The method of claim 1 wherein the data analysis is based at least in part on P-values.
3. The method of claim 1 wherein the data analysis is based at least in part on predictive determinism of variables and/or variable combinations from the data set.
4. The method of claim 1 further comprising the computer system performing the step of excluding statistical outliers from the data analysis.
5. The method of claim 1 further comprising the computer system performing the step of excluding data based on reduction of overall noise.
6. The method of claim 1 wherein the data analysis is based at least in part on predictive determinism of variables and/or variable combinations for predicting an outcome variable.
7. The method of claim 6 wherein the outcome variable is a variable from the data set.
8. The method of claim 6 wherein the outcome variable is calculated from variables from the data set.
9. The method of claim 6 further comprising the computer system performing the step of testing candidate outcome variables to select the outcome variable.
10. The method of claim 9 wherein the outcome variable is the candidate outcome variable with the highest signal to noise ratio.
11. The method of claim 6 further comprising a human performing the step of selecting the outcome variable.
12. The method of claim 6 further comprising a human performing the step of selecting candidate outcome variables, and the computer system performing the step of testing the candidate outcome variables to select the outcome variable.
13. The method of claim 1 wherein the data analysis analyzes combinations of variables and is capable of detecting relationship based on a combination of variables, where analysis of individual variables would not have detected such relationship.
14. The method of claim 1 wherein the data analysis is based at least in part on a local search.

15. The method of claim 1 wherein the data being analyzed evolves as the data analysis proceeds.
16. The method of claim 1 further comprising the computer system performing the steps of:
  - presenting the potentially valuable patterns to statistically untrained humans for feedback; and
  - receiving feedback from the statistically untrained humans.
17. A method implemented by a computer system for automatically reporting actionable insights in a data set, the method comprising the computer system performing the steps of:
  - analyzing data within the data set;
  - identifying potentially valuable patterns based on the data analysis;
  - developing actionable insights based on the identified potentially valuable patterns;
  - and
  - automatically generating reports that explain the actionable insights to a human reviewing the reports.
18. The method of claim 17 wherein, in the reports, a statistical importance of an actionable insight is represented by a numerical metric.
19. The method of claim 17 wherein, in the reports, statistical significance is a component in determining importance of an actionable insight.
20. The method of claim 17 wherein, in the reports, predictive determinism is a component in determining importance of an actionable insight.
21. The method of claim 17 wherein, in the reports, P-value is a component in determining importance of an actionable insight.
22. The method of claim 17 wherein, in the reports, an importance of an actionable insight is represented by a relative importance compared to the most important actionable insight.
23. The method of claim 17 wherein the computer system further performs the step of automatically generating an animation that explains the actionable insights to a human viewing the animation.
24. The method of claim 23 wherein the animation includes examples and/or graphs.
25. The method of claim 23 wherein the animation does not explain the data analysis to the human viewing the animation.
26. A combined computer/human method for automatically identifying actionable insights in a data set, the method comprising:

a computer system automatically and iteratively performing the steps of:  
analyzing data within the data set;  
identifying potentially valuable patterns based on the data analysis;  
presenting the potentially valuable patterns to statistically untrained humans  
for feedback; and  
receiving feedback from the statistically untrained humans; and  
iteratively developing actionable insights based on the steps taken by the computer  
system.

27. The method of claim 26 wherein the step of iteratively developing actionable insights is performed by the computer system.
28. The method of claim 26 wherein the step of iteratively developing actionable insights is not performed by statistically trained humans.
29. The method of claim 26 wherein the method is for identifying actionable insights for an entity, and the statistically untrained humans include customers of the entity.
30. The method of claim 26 wherein the method is for identifying actionable insights for an entity, and the statistically untrained humans include employees of the entity.
31. The method of claim 26 wherein the step of receiving feedback from the statistically untrained humans comprises receiving feedback from computer systems programmed by the statistically untrained humans.
32. The method of claim 26 wherein the method is for identifying actionable insights to improve a process.
33. The method of claim 26 wherein the method is for identifying actionable insights to improve fraud analytics.
34. The method of claim 26 wherein the method is for identifying actionable insights to improve supply chain analytics.
35. The method of claim 26 wherein the method is for identifying actionable insights to improve customer behavior analytics.
36. The method of claim 26 wherein the method is for identifying actionable insights to improve social media analytics.
37. The method of claim 26 wherein the method is for identifying actionable insights to improve web interaction analytics.
38. The method of claim 26 wherein the step of iteratively developing actionable insights is based on statistical hypothesis tests.



39. The method of claim 26 wherein the step of iteratively developing actionable insights comprises automatically identifying possible meaningful patterns within the data set.
40. The method of claim 26 wherein the step of iteratively developing actionable insights comprises automatically identifying possible meaningful subsets of data within the data set.
41. The method of claim 40 wherein the step of automatically identifying possible meaningful subsets of data within the data set is based on a statistical analysis of those subsets which includes calculation of p-values of statistical hypotheses for those subsets.
42. The method of claim 40 wherein the step of automatically identifying possible meaningful subsets of data within the data set is based on variance in the means of subsets.
43. The method of claim 40 wherein the step of automatically identifying possible meaningful subsets of data within the data set comprises iteratively refining possible meaningful subsets based on a local search method.
44. The method of claim 26 further comprising the computer system automatically performing the step of creating reports based on the data analysis.
45. The method of claim 26 wherein the step of iteratively developing actionable insights is based at least in part on predictive determinism of variables and/or variable combinations from the data set.
46. The method of claim 26 wherein the step of iteratively developing actionable insights comprises testing against a null hypothesis that all operators behave the same.
47. The method of claim 26 wherein the step of presenting the potentially valuable patterns to the humans comprises presenting the potentially valuable patterns to statistically untrained humans for voting on the possible hypotheses.
48. The method of claim 26 wherein the step of presenting the potentially valuable patterns to the humans comprises presenting the potentially valuable patterns to statistically untrained humans for tagging with comments.
49. The method of claim 26 wherein the step of presenting the potentially valuable patterns to the humans comprises presenting the potentially valuable patterns to statistically untrained humans for the suggestion of new possible hypotheses for actionable insights.
50. The method of claim 26 wherein the step of presenting the potentially valuable patterns to the humans comprises presenting the potentially valuable patterns to statistically untrained humans for the suggestion of refinements to possible hypotheses for actionable insights.
51. The method of claim 26 wherein the step of automatically identifying potentially valuable patterns comprises automatically identifying possible meaningful subsets of data

within the data set, and the step of presenting the potentially valuable patterns to the humans comprises presenting the potentially valuable patterns to statistically untrained humans for the suggestion of new possible meaningful subsets.

52. The method of claim 26 wherein the step of automatically identifying potentially valuable patterns comprises automatically identifying possible meaningful subsets of data within the data set, and the step of presenting the potentially valuable patterns to the humans comprises presenting the potentially valuable patterns to statistically untrained humans for the suggestion of refinements to the automatically identified possible meaningful subsets.

53. The method of claim 26 wherein the step of automatically identifying potentially valuable patterns comprises the computer system automatically identifying candidate variables or combination of variables for analysis, and the feedback received from the humans comprises suggestions for refinements to the candidate variables or combinations of variables.

54. The method of claim 26 wherein the step of presenting the potentially valuable patterns to the humans comprises presenting the potentially valuable patterns to statistically untrained humans for the suggestion of additional variables to be included in the steps performed by the computer system.

55. The method of claim 26 wherein the step of presenting the potentially valuable patterns to the humans comprises presenting the potentially valuable patterns to statistically untrained humans for the suggestion of additional data fields to be included in the steps performed by the computer system.

56. The method of claim 26 wherein the step of presenting the potentially valuable patterns to the humans comprises presenting the potentially valuable patterns to statistically untrained humans for the suggestion of types of regression to be included in the steps performed by the computer system.

57. The method of claim 26 wherein the step of presenting the potentially valuable patterns to the humans comprises presenting the potentially valuable patterns to statistically untrained humans for the suggestion of experiments to gather additional data.

58. The method of claim 26 wherein the human feedback includes structured feedback directly useable by the computer system.

59. The method of claim 58 wherein the structured feedback comprises words selected from a limited group of words.

60. The method of claim 58 wherein the structured feedback comprises sentences constructed from a limited group of words.

61. The method of claim 58 wherein the structured feedback comprises searches proposed by the statistically untrained humans.
62. The method of claim 26 wherein the human feedback includes unstructured feedback.
63. The method of claim 26 wherein the step of presenting the potentially valuable patterns to the humans includes presenting the feedback process as a game.
64. The method of claim 26 wherein the step of presenting the potentially valuable patterns to the humans includes an incentive for the humans to provide feedback.
65. The method of claim 64 wherein the incentive is based on uniqueness of the feedback.
66. The method of claim 64 wherein the incentive is based on effectiveness of the feedback.
67. The method of claim 26 further comprising the computer system performing the step of:  
automatically analyzing the received human feedback for patterns to improve the computer system's process for automatically identifying actionable insights.
68. The method of claim 26 wherein the step of presenting the potentially valuable patterns to humans for feedback comprises presenting different data to different humans.
69. The method of claim 26 wherein the feedback from humans comprises tagging of the potentially valuable patterns.
70. The method of claim 26 wherein the feedback from one human is available to another human when the other human is providing feedback.
71. The method of claim 26 further comprising the computer system automatically performing the step of applying a learning algorithm to the received human feedback.
72. The method of claim 26 further comprising the computer system automatically performing the step of detecting deterministic patterns in the received human feedback.
73. A computer program product for use with a computer, the computer program product comprising a tangible computer usable medium having a computer program code embodied therein for finding meaningful patterns in a data set, the computer program code automatically and iteratively performing the steps of:  
analyzing data within the data set;  
identifying potentially valuable patterns based on the data analysis;  
presenting the potentially valuable patterns to statistically untrained humans for feedback; and

receiving feedback from the statistically untrained humans, wherein actionable insights can be iteratively developed based on the steps performed by the computer program code.




FIG. 1

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 12/65921

## A. CLASSIFICATION OF SUBJECT MATTER

**IPC(8)** - G06F 17/30 (2013.01)**USPC** - 703/22

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

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8) Classification(s): G06F 17/30, G06F 19/00, H04N 7/32, H04N 7/30 (2013.01)

USPC Classification(s): 703/22, 705/346, 725/9

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

MicroPatent (US-G, US-A, EP-A, EP-B, WO, JP-bib, DE-C.B, DE-A, DE-T, DE-U, GB-A, FR-A); IEEE/IEEEExplore; Google/Google Scholar; IP.com; Search Terms Used: computer system, automatic, identify, actionable insight, analyze, data, human, people, intervention, predefine, rule, guideline, value, pattern, predict, hypothesis, deduct, interpretation, statistic, p value, noise reduction, outlier

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X -- Y	US 2009/0157419 A1 (BURSEY, B) June 18, 2009, paragraphs [0049], [0132], [0304]	1 ----- 2-25
Y	US 6799154 B1 (ARAGONES, JK et al.) September 28, 2004, column 4, lines 1-2, column 4, lines 25-46, column 5, lines 40-50,	2-4, 6-12, 15, 18-22
Y	WO 2010/071998 A1 (WONG, A et al.) July 1, 2010, page 10, lines 15-27, page 11, lines 3-8, page 14, lines 1-4	5, 10, 13
Y	US 2009/0094286 A1 (LEE HC, et al.) April 9, 2009, paragraphs [0051], [0056], [0083], [0084]	14, 17-25
Y	US 2002/0077787 A1 (RAPPAPORT, T et al.) June 20, 2002, abstract	16

☐ Further documents are listed in the continuation of Box C.

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&amp;" document member of the same patent family

Date of the actual completion of the international search

17 March 2013 (17.03.2013)

Date of mailing of the international search report

02 APR 2013

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents

P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Shane Thomas

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774

# INTERNATIONAL SEARCH REPORT

International application No.

PCT/US12/65921

## Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:  
because they relate to subject matter not required to be searched by this Authority, namely:
  
2. ☐ Claims Nos.:  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
  
3. ☐ Claims Nos.:  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

## Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

Group I: Claims 1-25; Group II: Claims 26-73

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fees must be paid.

Group I: Claims 1-25 are directed toward a method implemented by a computer system for automatically identifying actionable insights in a data set, the method comprising the computer system performing the steps of: analyzing data within the data set, wherein such analysis occurs without human intervention and without use of predefined rules; identifying potentially valuable patterns based on the data analysis; and developing actionable insights based on the identified potentially valuable patterns.

-"-Continued Within the Extra Page-" \*

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:  
Group I: Claims 1-25

### Remark on Protest



The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.



The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.



No protest accompanied the payment of additional search fees.

-Continued from Box No. III: Observations where unity of invention is lacking-

Group II: Claims 26-73 are directed toward a combined computer/human method for automatically identifying actionable insights in a data set, the method comprising: a computer system automatically and iteratively performing the steps of: analyzing data within the data set; identifying potentially valuable patterns based on the data analysis; presenting the potentially valuable patterns to statistically untrained humans for feedback; and receiving feedback from the statistically untrained humans; and iteratively developing actionable insights based on the steps taken by the computer system.

The common technical features among Groups I and II are a system for automatically identifying actionable insights in a data set, the method comprising the computer system performing the steps of: analyzing data within the data set, identifying potentially valuable patterns based on the data analysis; and developing actionable insights based on the identified potentially valuable patterns. However, these common technical features are disclosed by US 2009/0157419 A 1 to Bursey, hereafter referred to as Bursey. Bursey discloses a system for automatically identifying actionable insights in a data set (a Complex Event Processor is used to automatically provide actions that need to be taken, paragraph [0304]), the method comprising the computer system performing the steps of: analyzing data within the data set (the CEP identifies patterns in cloud data, paragraph [0304]), identifying potentially valuable patterns based on the data analysis (the CEP identifies patterns in cloud data, paragraph [0304]); and developing actionable insights based on the identified potentially valuable patterns (based on the identified pattern actions that need to be taken are generated and executed, paragraph [0304]).

Since the common technical features are previously disclosed by the Bursey reference, these common features are not special and so Groups I-II lack unity.